J. Fluid Mech. (2019), *vol.* 870, *pp.* 988–1036. © Cambridge University Press 2019
doi:10.1017/jfm.2019.212

Multi-scale proper orthogonal decomposition of complex fluid flows

M. A. Mendez[1,†], M. Balabane[2] and J.-M. Buchlin[1]

[1]von Karman Institute for Fluid Dynamics, Environmental and Applied Fluid Dynamics Department, Rhode-St-Genèse, 1640, Belgium

[2]Laboratoire Analyse, Géométrie et Applications, Université Paris 13, Villetaneuse, 93430, France

(Received 25 April 2018; revised 23 December 2018; accepted 11 March 2019; first published online 15 May 2019)

Data-driven decompositions are becoming essential tools in fluid dynamics, allowing for tracking the evolution of coherent patterns in large datasets, and for constructing low-order models of complex phenomena. In this work, we analyse the main limits of two popular decompositions, namely the proper orthogonal decomposition (POD) and the dynamic mode decomposition (DMD), and we propose a novel decomposition which allows for enhanced feature detection capabilities. This novel decomposition is referred to as multi-scale proper orthogonal decomposition (mPOD) and combines multi-resolution analysis (MRA) with a standard POD. Using MRA, the mPOD splits the correlation matrix into the contribution of different scales, retaining non-overlapping portions of the correlation spectra; using the standard POD, the mPOD extracts the optimal basis from each scale. After introducing a matrix factorization framework for data-driven decompositions, the MRA is formulated via one- and two-dimensional filter banks for the dataset and the correlation matrix respectively. The validation of the mPOD, and a comparison with the discrete Fourier transform (DFT), DMD and POD are provided in three test cases. These include a synthetic test case, a numerical simulation of a nonlinear advection–diffusion problem and an experimental dataset obtained by the time-resolved particle image velocimetry (TR-PIV) of an impinging gas jet. For each of these examples, the decompositions are compared in terms of convergence, feature detection capabilities and time–frequency localization.

Key words: computational methods, low-dimensional models

1. Introduction and motivation

The analysis of fluid flows hinges upon the identification of coherent flow features out of seemingly chaotic data. These structures could produce dynamic force loads (Barrero-Gil, Pindado & Avila 2012; Rahmanian *et al.* 2014), unstable patterns in the evolution of hydrodynamic instabilities (Charru & de Forcrand-Millard 2009; Gudmundsson & Colonius 2011; Melnikov, Kreilos & Eckhardt 2014; Schmid & Brandt 2014), transition to turbulence (Hussain 1986; Grinstein & DeVore 1996),

† Email address for correspondence: mendez@vki.ac.be

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

enhancement of heat and mass transfer (Mladin & Zumbrunnen 2000; Kuhn, Kenjereš & von Rohr 2010), noise generation (del Castillo-Negrete, Spong & Hirshman 2008; Nagarajan *et al.* 2018) and more.

The development of data processing algorithms for detecting coherent features is continuously fostered by the availability of datasets with ever-growing resolution, and by the ongoing big data revolution which permeates any area of applied science. Identifying the relevant features from high-dimensional datasets is the purpose of data-driven decompositions, which lay the foundation of model order reduction (MOR), data compression, filtering and pattern recognition, and which are nowadays greatly enlarging the toolbox at the researcher's disposal. Recent reviews on data-driven decomposition with emphasis on fluid dynamics are proposed by Rowley & Dawson (2017) and Taira *et al.* (2017), while an extensive overview on the impact of the big data revolution in turbulence research is presented by Pollard *et al.* (2017), while Duriez, Brunton & Noack (2017) give an overview of machine learning methods for flow control.

Data-driven decompositions represent the data as a linear combination of basis elements, referred to as modes, having spatial and temporal structures with a certain energy contribution. These decompositions can be classified into two major classes: energy based and frequency based. This work presents a generalized formulation of both approaches, analyses their respective limits and proposes a novel hybrid decomposition referred to as multi-scale proper orthogonal decomposition (mPOD).

In the following subsections, we present a literature review on the energy-based and frequency-based data-driven decompositions (§ 1.1) and outline the recent developments in hybrid methods (§ 1.2). An overview of the main contribution of this study and the manuscript organization is presented in § 1.3.

1.1. *Energy optimality or spectral purity?*

The fundamental energy based decomposition is the proper orthogonal decomposition (POD). This decomposition originates from the Karhunen–Loéve (KL) theorem (Loève 1977), which formulates the optimal (in a least-square sense) representation of a stochastic process as a series of orthogonal functions. This decomposition has set the basis for data processing methods which are known, depending on the field, as principal component analysis (PCA), Hotelling transform, empirical orthogonal function (EOF) or proper orthogonal decomposition (POD). Literature reviews on these applications can be found in the work of Cordier & Bergmann (2013), and the monographs from Jackson (1991) and Holmes *et al.* (2012).

The POD for data-driven analysis of turbulent flows was introduced to the fluid mechanics community by Lumley (1967, 1970), who focused on the spatial structures of the POD modes and their link to coherent structures in turbulent flows. The link between spatial and temporal structures was investigated by Aubry (1991), who presented a variant of the decomposition, referred to as bi-orthogonal decomposition (BOD, Aubry, Guyonnet & Lima 1991), which focused on the temporal structures of the modes. This link was made more explicit by Sirovich (1987, 1989, 1991), who proposed a simple algorithm, known as 'snapshot POD', to compute the entire decomposition.

It is nevertheless only thanks to the link between the POD and the well-known singular value decomposition (SVD) that the calculation of the POD modes of a discrete dataset could be formulated as a simple matrix factorization (see Kunisch & Volkwein 1999; Chatterjee 2000; Fahl 2000), in contrast to formulations based on

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

statistics or dynamical system theory (such as Berkooz 1992; Berkooz, Holmes & Lumley 1993; Holmes *et al.* 1997; Noack *et al.* 2003). For an introduction to the SVD and a historical account on its development, the reader is referred to Stewart (1993) and Golub (2013).

This simple matrix factorization framework has made the POD an extremely popular tool for both experimental and numerical fluid mechanics. Typical applications include the identification of coherent structures from experimental data (e.g. Schrijer, Sciacchitano & Scarano 2014; Mallor *et al.* 2018), video analysis for adaptive masking or image pre-processing (Mendez *et al.* 2017; Mendez, Scelzo & Buchlin 2018*b*), flow sensing (Venturi & Karniadakis 2004; Willcox 2006), flow control (Bergmann & Cordier 2008; Brunton & Noack 2015), reduced-order modelling (Deane *et al.* 1991), experimental validation of computational fluid dynamics data (Blanchard *et al.* 2014), data-driven identification of nonlinear systems (Brunton, Proctor & Kutz 2016*b*; Loiseau, Noack & Brunton 2018) and more (see Pollard *et al.* 2017). Moreover, extended versions of the decomposition, in which the dataset to decompose is constructed by assembling different quantities (e.g. velocities and concentrations), have also found application in the correlation analysis of coupled phenomena (Glezer, Kadioglu & Pearlstein 1989; Maurel, Borée & Lumley 2001; Borée 2003; Duwig & Iudiciani 2009; Antoranz *et al.* 2018).

While the POD optimality guarantees the identification of the most energetic contributions within a few modes, the interpretation of these taken individually, and thus their usage for feature detection, can become difficult. As later illustrated with exemplary test cases, there exist situations in which different phenomena, possibly occurring at largely different frequencies, have very similar energy content. In such cases, an energy-based formulation cannot distinguish the various contributions, which consequently share the same POD structures. Frequency-based methods, conversely, assign to each mode a single frequency. Such methods are data-driven adaptations of the discrete Fourier transform (DFT) and infer the frequency of each mode from the data rather than defining it *a priori* (like in the DFT) from the time discretization.

The fundamental frequency-based and data-driven method is the dynamic mode decomposition (DMD). This decomposition originates from the Koopman theory (Koopman 1931), which consists in studying a nonlinear system by mapping it onto a linear one of much larger size (generally infinite-dimensional). This theory was then used by Mezić (2005) to show how finite-dimensional approximations of such a Koopman (linear) dynamical system could be used for MOR applications. The extraction of such a finite approximation from the data was proposed by Rowley *et al.* (2009) and Schmid (2010). An introduction to the Koopman framework is proposed by Brunton *et al.* (2016*a*), while extensive reviews on its connection to the DMD are proposed by Budišić, Mohr & Mezić (2012), Mezić (2013) and Tu *et al.* (2014). After the first two formulations, briefly reviewed in this work, many extensions of the algorithm have been developed; examples are the optimized DMD (Chen, Tu & Rowley 2012), the sparsity-promoting DMD (Jovanović, Schmid & Nichols 2014), the randomized DMD (Erichson, Brunton & Kutz 2017) or the multi-resolution DMD (Kutz, Fu & Brunton 2016*b*). An overview of the DMD and its application is proposed by Kutz *et al.* (2016*a*).

Besides the data-driven selection of the frequency of each mode, the advantage of the DMD over the DFT is to allow these to grow or decay exponentially, being each frequency complex. This makes the DMD suited for data-driven stability analysis (Schmid & Brandt 2014) and feedback control design for linear systems (Rowley & Dawson 2017).

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

The constraint of fully harmonic modes, on the other hand, can pose significant challenges to frequency-based decompositions for nonlinear datasets. A purely harmonic decomposition cannot easily represent frequency modulations, frequency and phase jitter or temporally localized events. A large number of modes required for the harmonic description of these phenomena in one domain (e.g. time) results in a substantial redundancy in the structures in the other domain (e.g. space). Moreover, since harmonics have infinite temporal support (they do not 'start' nor 'finish') their localization capabilities are particularly weak. Extending the frequency domain to the complex plane, as in the DMD, further amplifies these problems in the presence of nonlinear phenomena such as saturation of growth or decay rates.

The limits of harmonic bases have motivated the development of time–frequency analysis and wavelet theory (Mallat 2009; Kaiser 2010), which has found many applications in the multi-scale modal analysis for fluid flows (Farge 1992; Berkooz *et al.* 1994; De Souza *et al.* 1999; Rinoshika & Zhou 2005; Schneider & Vasilyev 2010). Wavelet decompositions are nevertheless not data-driven decomposition since one defines *a priori* the set of basis elements which have a given scale (frequency) and extension (duration). As described in § 3 and appendix A, the proposed mPOD borrows several ideas from the wavelet literature and combines them with the optimality constraints of the POD.

1.2. *The need for hybrid decompositions*

Focusing on data-driven methods, the limits of energy based (POD) or frequency based (DMD) approaches have been recently debated in the literature, and Noack (2016) has recently discussed the quest for hybrid decompositions. The first single-harmonic variant of the POD was proposed by Lumley (1970) and formulated as a natural extension of the POD for stationary and homogeneous flows (see Glauser, Leib & George 1987; George 1988). This approach consists in computing the POD structures as eigenvectors of the cross-spectral density matrix (see Citriniti & George 2000; Gordeyev & Thomas 2000; Gudmundsson & Colonius 2011) and is referred to by several authors (e.g. Picard & Delville 2000; Taira *et al.* 2017; Towne, Schmidt & Colonius 2018) as spectral proper orthogonal decomposition (SPOD). Bourgeois, Noack & Martinuzzi (2013) uses a pre-filtering of the data with a low-pass Morlet filter before computing the POD, to limit the frequency content of the resulting modes. Cammilleri *et al.* (2013) proposes a combination of POD and DMD, referred to as Cronos–Koopman analysis, which consists in performing DMD on the temporal basis of the POD, to force purely harmonic temporal modes. Noack *et al.* (2016) uses a recursive approach for forcing the orthogonality in the DMD modes while minimizing the loss of spectral purity.

It is therefore evident that both the energy maximization and the spectral purity are too restrictive constraints, and an ideal decomposition should match the two approaches, possibly allowing for switching from one to the other. A decisive step towards this direction was achieved by Sieber, Paschereit & Oberleithner (2016), who proposed a novel decomposition blending the POD and the DFT. This decomposition was also named spectral proper orthogonal decomposition (SPOD), although it has nothing to do with the frequency-based POD that is also referred to as SPOD by Picard & Delville (2000), Taira *et al.* (2017) and Towne *et al.* (2018).

Sieber's SPOD consists of using a low-pass filter along all the diagonals of the correlation matrix, before proceeding the Sirovich's 'snapshot POD'. The idea of such diagonal filtering arises from Szegö theorem (Grenander 2001; Gray 2005)

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

which states that the eigenbasis of a Toeplitz circulant matrix is a Fourier basis. Therefore, the more the correlation matrix resembles a circulant Toeplitz matrix, the more a POD resembles a DFT. A correlation matrix, however, approaches a Toeplitz circulant form only for a stationary process (Brockwell & Davis 1987; Gray 2005) and the diagonal filter used in the SPOD artificially forces such a pattern, thus forcing the resulting eigenvectors to approach the DFT basis. Depending on the filter size, one moves from a DFT (very large filter kernel) to a POD (very short filter kernel). Between these two limits, a good compromise between energy optimality and spectral purity can be obtained.

The main limitation of this method, however, is that this filtering operation can significantly alter the correlation matrix of non-stationary processes, compromising its symmetry and thus the orthogonality of its eigenvectors. Such operation can be seen as a smoothing of the nonlinearities (Sieber *et al.* 2016; Sieber, Paschereit & Oberleithner 2017), but the decision on the filter size is not supported by a mathematical link between the diagonal filtering and the spectral content of the resulting eigenvectors.

1.3. *Contribution of this work and manuscript organization*

Section 2 presents a general matrix factorization that is common to all the decompositions, and sets the POD, DFT and DMD in this generalized form. This framework is then used in § 3 to derive the novel multi-scale proper orthogonal decomposition (mPOD).

The mPOD proposed in this work combines ideas from the windowed SPOD in Sieber *et al.* (2016) and the SPOD in Towne *et al.* (2018). Instead of filtering the correlation matrix, as in Sieber *et al.* (2016), the mPOD decomposes it into the contributions of different scales using the multi-resolution (MRA) architecture from wavelet theory. Instead of computing various eigenbases from the spectra of different portions of the data, as in Towne *et al.* (2018), the mPOD computes eigenbases on the correlation matrix of different scales. The derivation of the mPOD in § 3 is carried out by analysing the impact of the filtering process on the spectra of the correlation matrix and the POD modes. Finally, the mPOD is tested and compared to POD, DFT and DMD on an illustrative synthetic test case (§ 4), a numerical test case (§ 5) and an experimental test case (§ 6). Conclusions and perspectives for future applications are presented in § 7.

2. Data decompositions as matrix factorizations

Let $\boldsymbol{D}(\boldsymbol{x_i}, t_k) = \boldsymbol{D}[\boldsymbol{i}, k] \in \mathbb{R}^{n_s \times n_t}$ be a matrix collecting the set of time realizations (snapshot) of a real variable. In the time domain, the data are sampled with a uniform temporal discretization $\{t_k = (k-1)\Delta t\}_{k=1}^{n_t}$. In the spatial domain, the data are sampled on a Cartesian grid $\boldsymbol{x_i} \in \mathbb{R}^{n_x \times n_y}$, with $\boldsymbol{i} \in [1, \ldots, n_s]$ a matrix linear index, $n_s = n_C n_x n_y$ the number of points and n_C the number of components of the data in case of vector quantities (e.g. $n_C = 1$ for a pressure field, $n_C = 3$ for a three-dimensional velocity field). After reshaping each snapshot of the data into a column vector $\boldsymbol{d}_k[\boldsymbol{i}] \in \mathbb{R}^{n_s \times 1}$, the data matrix read

$$\boldsymbol{D}[\boldsymbol{i}, k] = \begin{bmatrix} \boldsymbol{d}_1[1] & \cdots & \boldsymbol{d}_k[1] & \cdots & \boldsymbol{d}_{n_t}[1] \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \boldsymbol{d}_1[n_s] & \cdots & \boldsymbol{d}_k[n_s] & \cdots & \boldsymbol{d}_{n_t}[n_s] \end{bmatrix}. \tag{2.1}$$

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

The criterion followed to reshape each snapshot into a vector is irrelevant, provided that the same is used when reshaping back the results of the decomposition. The scope of any discrete decomposition is to break this matrix into the sum of rank-1 contributions, referred to as modes, written in variable separated form. Each mode has a spatial structure $\boldsymbol{\phi}_r[\boldsymbol{i}]$, a temporal structure $\boldsymbol{\psi}_r[k]$ and an amplitude σ_r:

$$\boldsymbol{D}[\boldsymbol{i}, k] = \sum_{r=1}^{rk(\boldsymbol{D})} \sigma_r \boldsymbol{\phi}_r[\boldsymbol{i}] \boldsymbol{\psi}_r[k]. \tag{2.2}$$

A data matrix of rank $rk(\boldsymbol{D}) = \min(n_t, n_s)$ is said to have a full rank, while truncating the summation at $r_c < rk(\boldsymbol{D})$ leads to an approximation $\tilde{\boldsymbol{D}}$ of rank r_c. In the MOR community, this approximation is referred to as a reduced-order model of the data.

In order to let the scalars σ_r be fully representative of the energy of each mode, both the spatial and the temporal structures must have unitary energy (norm). The notion of energy is defined by the choice of inner product in the space and in the time domain. Although different definitions have been proposed in the literature (Holmes *et al.* 2012; Cordier & Bergmann 2013), we herein focus on the classical L^2 inner product in its continuous and discrete forms. We shall moreover assume that any data realization (and any spatial or temporal structure) is the discrete representation (sampling) of a continuous and square integrable function. Therefore, let the vector $\boldsymbol{\phi}_r[\boldsymbol{i}] \in \mathbb{C}^{n_s \times 1}$ collect n_s samples of its continuous counterpart $\boldsymbol{\phi}_r(\boldsymbol{x})$, over a spatial domain $x \in \Omega$ discretized with a Cartesian grid $\boldsymbol{x}_i$. An estimation of the average energy content of this function is provided by the Euclidean inner product

$$\|\boldsymbol{\phi}_r(\boldsymbol{x})\|_2^2 = \frac{1}{\Omega} \int_\Omega \boldsymbol{\phi}_r(\boldsymbol{x}) \overline{\boldsymbol{\phi}}_r(\boldsymbol{x}) \, \mathrm{d}\Omega \approx \frac{1}{n_s} \sum_{\mathrm{i}=1}^{n_s} \boldsymbol{\phi}_r[\boldsymbol{i}] \overline{\boldsymbol{\phi}}_r[\boldsymbol{i}] = \frac{1}{n_s} \langle \boldsymbol{\phi}_r, \boldsymbol{\phi}_r \rangle = \frac{1}{n_s} \boldsymbol{\phi}_r^\dagger \boldsymbol{\phi}_r, \tag{2.3}$$

where the spatial structure is also reshaped as a column vector, the overline denotes complex conjugation, $\langle \boldsymbol{a}, \boldsymbol{b} \rangle = \boldsymbol{a}^\dagger \boldsymbol{b}$ is the Euclidean inner product of two vectors $\boldsymbol{a}, \boldsymbol{b}$ and † represents the Hermitian transpose. Similarly, assuming that the vector $\boldsymbol{\psi}_r[k] \in \mathbb{C}^{n_t \times 1}$ collects n_t samples of the function $\boldsymbol{\psi}_r(t)$, defined in a time domain $t \in [0, T]$, the averaged energy content in the time domain reads:

$$\|\boldsymbol{\psi}_r(t)\|_2^2 = \frac{1}{T} \int_T \boldsymbol{\psi}_r(t) \overline{\boldsymbol{\psi}}_r(t) \, \mathrm{d}T \approx \frac{1}{n_t} \sum_{k=1}^{n_t} \boldsymbol{\psi}_r[k] \overline{\boldsymbol{\psi}}_r[k] = \frac{1}{n_t} \langle \boldsymbol{\psi}_r, \boldsymbol{\psi}_r \rangle = \frac{1}{n_t} \boldsymbol{\psi}_r^\dagger \boldsymbol{\psi}_r. \tag{2.4}$$

Equations (2.3) and (2.4) are numerical approximations of the continuous inner products using the right endpoint rule, and are valuable only if both the spatial and the temporal discretizations are Cartesian. More advanced quadrature methods or non-uniform discretization requires the definition of a weighted inner product (Gordeyev & Thomas 2000; Cordier & Bergmann 2013; Volkwein 2013). To simplify the following derivations, we here focus on Euclidean inner products and uniformly sampled data.

From (2.3)–(2.4), it is clear that the normalization $\|\boldsymbol{\phi}_r\|_2 = \|\boldsymbol{\psi}_r\|_2 = 1$, at a discrete level, implies that the amplitudes σ_r in (2.2) must be normalized by $\sqrt{n_s n_t}$ to produce a grid-independent estimation of the energy of each mode. Before discussing further the link between the energy contributions σ_r and the energy of the entire dataset, it is useful to arrange the spatial and the temporal structures into matrices

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

$\boldsymbol{\Phi} = [\boldsymbol{\phi}_1[\boldsymbol{i}], \ldots, \boldsymbol{\phi}_{rk(\boldsymbol{D})}[\boldsymbol{i}]] \in \mathbb{C}^{n_s \times rk(\boldsymbol{D})}$, $\boldsymbol{\Psi} = [\boldsymbol{\psi}_1[k], \ldots, \boldsymbol{\psi}_{rk(\boldsymbol{D})}[k]] \in \mathbb{C}^{n_t \times rk(\boldsymbol{D})}$. Any decomposition of the form (2.2) becomes a matrix factorization of the form

$$\boldsymbol{D} = \sum_{r=1}^{rk(\boldsymbol{D})} \sigma_r \boldsymbol{\phi}_r \boldsymbol{\psi}_r^{\mathrm{T}} = \boldsymbol{\Phi}\boldsymbol{\Sigma}\boldsymbol{\Psi}^{\mathrm{T}}, \tag{2.5}$$

where $\boldsymbol{\Sigma} = \mathrm{diag}[\sigma_1, \sigma_2, \ldots, \sigma_{rk(\boldsymbol{D})}]$ is the diagonal matrix containing the energy contribution of each mode, and the superscript $^{\mathrm{T}}$ denotes matrix transposition. Besides allowing for a compact notation, (2.5) allows for visualizing the operation as a projection of the dataset onto a spatial basis $\boldsymbol{\Phi}$ and a temporal basis $\boldsymbol{\Psi}$; the first is a base for the columns of $\boldsymbol{D}$; the second is a base for its rows.

It is worth noticing that, since $\boldsymbol{\Sigma}$ is diagonal and $\boldsymbol{\Phi}$ and $\boldsymbol{\Psi}$ have normalized columns, the factorization in (2.5) is entirely defined by either $\boldsymbol{\Phi}$ or $\boldsymbol{\Psi}$. In this work, we focus on projections in the time (row) domain of the dataset, i.e. (2.5) is defined by the temporal structures $\boldsymbol{\Psi}$. From (2.5), a matrix inversion and a column-wise normalization allow for computing the spatial structures and the amplitudes:

$$\boldsymbol{\Phi}\boldsymbol{\Sigma} = \boldsymbol{D}(\boldsymbol{\Psi}^{\mathrm{T}})^{-1} = \boldsymbol{C} \rightarrow \boldsymbol{\Phi} = \boldsymbol{C}\boldsymbol{\Sigma}^{-1}, \tag{2.6}$$

where $\sigma_r = \|\boldsymbol{C}_r\|_2$ and $\boldsymbol{C}_r$ is r^{th} column of the matrix $\boldsymbol{C} = \boldsymbol{D}(\boldsymbol{\Psi}^{\mathrm{T}})^{-1}$.

Particularly convenient is the case of an orthonormal temporal basis, for which $\boldsymbol{\Psi}^{\dagger}\boldsymbol{\Psi} = \boldsymbol{\Psi}^{\mathrm{T}}\overline{\boldsymbol{\Psi}} = \boldsymbol{I}$, with $\boldsymbol{I}$ the identity matrix. In this case, the inversion becomes $(\boldsymbol{\Psi}^{\mathrm{T}})^{-1} = \overline{\boldsymbol{\Psi}}$ and the energy contributions, which become $\sigma_r = \|\boldsymbol{D}\boldsymbol{\psi}_r\|_2 = \|\sigma_r\boldsymbol{\phi}_r\|_2$, can be computed independently for each mode.

Among the infinite possible bases $\boldsymbol{\Psi}$, we focus on the three most popular ones: the POD, the DFT and the DMD. Their matrix form is discussed in the following subsections, where the subscripts $\mathcal{P}$, $\mathcal{F}$ and $\mathcal{D}$ are used to distinguish the decompositions.

Before proceeding, it is worth commenting on two important aspects of any decompositions: (i) the treatment of the time (column-wise) average, (ii) the link between the contributions σ_r and the total energy content in $\boldsymbol{D}$. Concerning the first, one can see from (2.2) and (2.5) that removing the time average is equivalent to taking the normalized constant vector $\boldsymbol{\psi}_\mu = \underline{1}/\sqrt{n_t} \in \mathbb{R}^{n_t \times 1}$ as one of the temporal structures, such that

$$\boldsymbol{D}[\boldsymbol{i}, k] = d_\mu[\boldsymbol{i}]\boldsymbol{\psi}_\mu[k] + \breve{\boldsymbol{D}}[\boldsymbol{i}, k] = \sigma_\mu\boldsymbol{\phi}_\mu[\boldsymbol{i}]\boldsymbol{\psi}_\mu[k] + \sum_{r=1}^{rk(\boldsymbol{D})-1} \breve{\sigma}_r\breve{\boldsymbol{\phi}}_r[\boldsymbol{i}]\breve{\boldsymbol{\psi}}_r[k], \tag{2.7}$$

where $d_\mu[\boldsymbol{i}] = \boldsymbol{D}\boldsymbol{\psi}_\mu$ is proportional to the average column of the dataset, and the energy content $\sigma_\mu = \|d_\mu\|$ is proportional to the averaged correlation level, to be normalized by $\sqrt{n_s n_t}$ to provide a grid independent estimation. The $\breve{\boldsymbol{D}}[\boldsymbol{i}, k]$ is the zero-mean-shifted dataset and the last summation is its modal decomposition (distinguished by an accent $\breve{\bullet}$). The zero-mean-shifted dataset has $rk(\breve{\boldsymbol{D}}) = n_t - 1$ and all its temporal structures – if an orthogonal decomposition is considered – must have zero mean. This is naturally the case for the DFT, but not for the POD or the DMD.

Concerning the second point, the energy content of the dataset can be defined via different matrix norms. Using the most classical Frobenius norm yields

$$\|\boldsymbol{D}\|_F^2 = \mathrm{Tr}(\underbrace{\boldsymbol{D}^{\dagger}\boldsymbol{D}}_{\boldsymbol{K}}) = \mathrm{Tr}(\underbrace{\boldsymbol{D}\boldsymbol{D}^{\dagger}}_{\boldsymbol{C}}) = \sum_{r=1}^{rk(\boldsymbol{D})} \lambda_r, \tag{2.8}$$

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

where 'Tr' indicates the trace of the matrices $\boldsymbol{K} = \boldsymbol{D}^\dagger \boldsymbol{D} \in \mathbb{R}^{n_t \times n_t}$ and $\boldsymbol{C} = \boldsymbol{D}\boldsymbol{D}^\dagger \in \mathbb{R}^{n_s \times n_s}$. These are the finite-dimensional estimators of the two-point temporal and spatial correlation tensors (Holmes *et al.* 2012; Cordier & Bergmann 2013) respectively, using the inner products in (2.3)–(2.4). These symmetric positive definite matrices share the same non-zero eigenvalues λ_r. Introducing (2.5) in (2.8) gives two ways of measuring the total energy from the contribution of each mode:

$$\|\boldsymbol{D}\|_F^2 = \mathrm{Tr}(\underbrace{\boldsymbol{\Psi}\boldsymbol{\Sigma}\boldsymbol{\Phi}^\dagger\boldsymbol{\Phi}\boldsymbol{\Sigma}\boldsymbol{\Psi}^\dagger}_{K}) = \mathrm{Tr}(\underbrace{\boldsymbol{\Phi}\boldsymbol{\Sigma}\boldsymbol{\Psi}^\dagger\boldsymbol{\Psi}\boldsymbol{\Sigma}\boldsymbol{\Phi}^\dagger}_{C}) = \sum_{r=1}^{rk(\boldsymbol{D})} \lambda_r. \tag{2.9}$$

It is therefore evident that it is not possible to compute the total energy of the data by solely using the energy contributions of the modes, unless both the spatial and the temporal structures are orthonormal ($\boldsymbol{\Phi}^\dagger\boldsymbol{\Phi} = \boldsymbol{\Psi}^\dagger\boldsymbol{\Psi} = \boldsymbol{I}$). This case is described in § 2.1.

2.1. *Matrix form of the proper orthogonal decomposition (POD)*

The POD temporal structures $\boldsymbol{\Psi}_\mathcal{P}$ are the eigenvectors of the temporal correlation matrix $\boldsymbol{K} = \boldsymbol{D}^\dagger\boldsymbol{D}$. If $\boldsymbol{D}$ is real, as in all the applications considered in this work, these are orthonormal and real vectors:

$$\boldsymbol{K} = \boldsymbol{\Psi}_\mathcal{P}\boldsymbol{\Lambda}_\mathcal{P}\boldsymbol{\Psi}_\mathcal{P}^\mathrm{T} = \sum_{r=1}^{rk(\boldsymbol{D})} \lambda_{\mathcal{P}_r}\boldsymbol{\psi}_{\mathcal{P}_r}\boldsymbol{\psi}_{\mathcal{P}_r}^\mathrm{T}. \tag{2.10}$$

Introducing this definition in (2.9) is particularly revealing: a comparison of (2.10) to the first factorization in (2.9) yields $\boldsymbol{\Lambda}_\mathcal{P} = \boldsymbol{\Sigma}_P\boldsymbol{\Phi}_\mathcal{P}^\mathrm{T}\boldsymbol{\Phi}_\mathcal{P}\boldsymbol{\Sigma}_\mathcal{P}$. Since this eigenvalue matrix is diagonal, one recovers $\boldsymbol{\Phi}_\mathcal{P}^\mathrm{T}\boldsymbol{\Phi}_\mathcal{P} = \boldsymbol{I}$ and $\boldsymbol{\Lambda}_\mathcal{P} = \boldsymbol{\Sigma}_\mathcal{P}^2$. Therefore, also the spatial structures are real and orthonormal and thus eigenvectors of the spatial correlation matrix $\boldsymbol{C} = \boldsymbol{D}\boldsymbol{D}^\dagger = \boldsymbol{\Phi}_\mathcal{P}\boldsymbol{\Lambda}_\mathcal{P}\boldsymbol{\Phi}_\mathcal{P}^\mathrm{T}$. Therefore, the energy contribution of each mode becomes $\sigma_{\mathcal{P}_r} = \sqrt{\lambda_r}$, and the energy of the dataset can be recovered from $\boldsymbol{\Sigma}_\mathcal{P}$.

For the choice of temporal basis in (2.10), (2.5) yields the well-known singular value decomposition (SVD), here denoted as $\boldsymbol{D} = \boldsymbol{\Phi}_\mathcal{P}\boldsymbol{\Sigma}_\mathcal{P}\boldsymbol{\Psi}_\mathcal{P}^\mathrm{T}$. The optimality of the decomposition is thus guaranteed by the Eckart–Young–Mirsky theorem, which states that the rank $r_c < rk(\boldsymbol{D})$ approximation of a matrix $\boldsymbol{D}$ obtained from a truncated SVD minimizes the norm of the error

$$Err(r_c) = \left\| \boldsymbol{D} - \sum_{r=1}^{r_c} \sigma_{\mathcal{P}r}\boldsymbol{\phi}_{\mathcal{P}r}\boldsymbol{\psi}_{\mathcal{P}r}^\mathrm{T} \right\|_2 = \sigma_{\mathcal{P}r_c+1}. \tag{2.11}$$

Therefore, each coefficients $\sigma_{\mathcal{P}_r}$ represents also the L^2 error produced by a rank $r-1$ approximation, and the error decay $Err(r_c) \to 0$ as $r_c \to rk(\boldsymbol{D})$ is the strongest possible.

The POD from Lumley (1967, 1970) computes the spatial structures $\boldsymbol{\Phi}_\mathcal{P}$ from the eigendecomposition of $\boldsymbol{C}$; the BOD from Aubry (1991) computes the temporal structures (referred to as BOD modes) $\boldsymbol{\Psi}_\mathcal{P}$ from the eigendecomposition of $\boldsymbol{K}$. The POD from Sirovich (1987, 1989, 1991) computes both $\boldsymbol{\Phi}_\mathcal{P}$ and $\boldsymbol{\Psi}_\mathcal{P}$ by first solving (2.10) (less demanding if $n_t \ll n_s$) and then using (2.6) (which reads $\boldsymbol{\Phi}_\mathcal{P} = \boldsymbol{D}\boldsymbol{\Psi}_\mathcal{P}\boldsymbol{\Sigma}_\mathcal{P}^{-1}$).

Concerning the impact of the mean removal before computing the POD from (2.7), the POD is equipped with a vector of constant in its temporal basis only if such a vector is an eigenvector of $\boldsymbol{K}$. This implies:

$$\boldsymbol{K}\underline{1} = \sigma_\mu^2\underline{1} \quad \Longleftrightarrow \quad \sum_{j=1}^{n_t} \boldsymbol{K}[i,j] = \sigma_\mu^2 \quad \forall j \in [1, n_t] \tag{2.12}$$

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

that is the sum over any row of $\boldsymbol{K}$ is equal to the corresponding eigenvalue σ_μ^2. This occurs for a stationary process, for which the temporal correlation matrix tends towards a circulant form. For a non-stationary process, the mean removal imposes an additional constraint that alters the connection with the SVD. In such condition, one can still recover the result on optimality by considering the decomposition process as a linear affine transformation (see Miranda, Borgne & Le Bontempi 2008, for a derivation).

Finally, it is worth highlighting that the energy optimality is linked to the decomposition uniqueness: in the case of repeated singular values, which occurs when different coherent phenomena have a similar energy content (as shown in § 4), there exist infinite possible choices of singular vectors $\boldsymbol{\Phi}_\mathcal{P}$ and $\boldsymbol{\Psi}_\mathcal{P}$.

2.2. *Matrix form of the discrete Fourier transform (DFT)*

The DFT temporal structures $\boldsymbol{\Psi}_\mathcal{F}$ are defined *a priori*, regardless of the dataset at hand, as columns of the well-known Fourier matrix

$$\boldsymbol{\Psi}_\mathcal{F} = \frac{1}{\sqrt{n_t}} \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & w & w^2 & \cdots & w^{n_t-1} \\ \vdots & \vdots & \ddots & \vdots & \\ 1 & w^{n_t-1} & w^{2(n_t-1)} & \cdots & w^{(n_t-1)^2} \end{bmatrix}, \tag{2.13}$$

where $w = \exp(2\pi \mathrm{j}/n_t)$, with $\mathrm{j} = \sqrt{-1}$. Each temporal structure is thus a complex exponential $\boldsymbol{\psi}_r[k] = \exp(2\pi \mathrm{j} f_r(k-1)\Delta t) = \exp(2\pi \mathrm{j}(r-1)(k-1)/n_t)$ with real frequency discretization $f_r = (r-1)\Delta f$, with $\Delta f = 1/(n_t \Delta t) = f_s/n_t$ the frequency resolution and f_s the sampling frequency.

Besides being orthonormal ($\boldsymbol{\Psi}_\mathcal{F}^{-1} = \boldsymbol{\Psi}_\mathcal{F}^\dagger$), this Vandermonde matrix is also symmetric ($\boldsymbol{\Psi}_\mathcal{F}^\mathrm{T} = \boldsymbol{\Psi}_\mathcal{F}$). The resulting spatial structures $\boldsymbol{\Phi}_\mathcal{F} = \boldsymbol{D}\overline{\boldsymbol{\Psi}}_\mathcal{F}\boldsymbol{\Sigma}_\mathcal{F}^{-1}$ from (2.6) are complex and generally not orthogonal. Finally, since the DFT satisfies (2.7) by construction, the zero-mean shifting has no effect on the decomposition.

2.3. *Matrix form of the dynamic mode decomposition (DMD)*

The DMD temporal structures $\boldsymbol{\Psi}_\mathcal{D}$ are computed under the assumption that a linear dynamical system can well approximate the dataset. It is thus assumed that a propagator matrix $\boldsymbol{P} \in \mathbb{R}^{n_s \times n_s}$ maps each column of the data matrix $\boldsymbol{D}$ onto the following one via simple matrix multiplication

$$\boldsymbol{d}_k = \boldsymbol{P}\boldsymbol{d}_{k-1} = \boldsymbol{P}^{k-1}\boldsymbol{d}_1. \tag{2.14}$$

As each time step k involves the $k-1$ power of P acting on the initial data $\boldsymbol{d}_1$, the evolution of such linear system depends on the eigendecomposition $\boldsymbol{P} = \boldsymbol{S}\boldsymbol{\Lambda}\boldsymbol{S}^{-1}$ of the propagator, since $\boldsymbol{P}^{k-1} = \boldsymbol{S}\boldsymbol{\Lambda}^{k-1}\boldsymbol{S}^{-1}$:

$$\boldsymbol{d}_k[\boldsymbol{i}] = \boldsymbol{S}\boldsymbol{\Lambda}^{k-1}(\boldsymbol{S}^{-1}\boldsymbol{d}_1[\boldsymbol{i}]) = \boldsymbol{S}\boldsymbol{\Lambda}^{k-1}a_0 = \sum_{r=1}^{n_s} a_0[r]s_r[\boldsymbol{i}]\lambda_r^{k-1}. \tag{2.15}$$

The columns of the eigenvector matrix $\boldsymbol{S} = [s_1, s_2, \ldots, s_{n_s}]$ represent the spatial basis of such evolution. Observe that the possible growth/decay of each mode makes the notion of mode 'amplitude' particularly cumbersome in the DMD. Instead, the

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

vector $\boldsymbol{a}_0 = \boldsymbol{S}^{-1}\boldsymbol{d}_1$ is the projection of the initial data $\boldsymbol{d}_1$ onto the eigenvectors of the propagator. To arrange (2.15) into the factorization form (2.5), one should first arrange the vector $\boldsymbol{a}_0$ into a diagonal matrix $\boldsymbol{A}_0 = \mathrm{diag}(a_0)$ and then build the temporal basis from the powers of the eigenvalues. The resulting factorization can be written as

$$\boldsymbol{D} = \boldsymbol{S}\boldsymbol{A}_0\boldsymbol{Z}^{\mathrm{T}} \quad \text{with } \boldsymbol{Z}^{\mathrm{T}} = \begin{bmatrix} 1 & \lambda_1 & \lambda_1^2 & \cdots & \lambda_1^{n_t} \\ 1 & \lambda_2 & \lambda_2^2 & \cdots & \lambda_2^{n_t} \\ \vdots & \vdots & \ddots & \vdots & \\ 1 & \lambda_{n_s} & \lambda_{n_s}^2 & \cdots & \lambda_{n_s}^{n_t} \end{bmatrix}. \tag{2.16}$$

The temporal basis of the DMD is thus obtained by normalizing each column of $\boldsymbol{Z}$ independently, that is $\boldsymbol{\psi}_{\mathcal{D}r} = Z_r/\|Z_r\|$, and the spatial structures can be computed using the projection in (2.6) with no need for computing the eigenvectors $\boldsymbol{S}$. In this final step, it is important to observe that the DMD temporal basis is generally not orthogonal and thus the matrix inversion is generally unavoidable.

To compute the DMD, one needs to calculate the eigenvalues of the best (in a L^2 sense) propagator describing the dataset. This propagator can be easily defined by arranging the data matrix into two shifted portions, $\boldsymbol{D}_1 = [\boldsymbol{d}_1, \boldsymbol{d}_2, \ldots, \boldsymbol{d}_{n_t-1}]$ and $\boldsymbol{D}_2 = [\boldsymbol{d}_2, \boldsymbol{d}_3, \ldots, \boldsymbol{d}_{n_t}]$, containing $n_t - 1$ realization:

$$\boldsymbol{D}_2 \approx [\boldsymbol{P}\boldsymbol{d}_1, \boldsymbol{P}\boldsymbol{d}_2, \ldots, \boldsymbol{P}\boldsymbol{d}_{n_t-1}] = \boldsymbol{P}\boldsymbol{D}_1 \Longleftrightarrow \boldsymbol{P} = \boldsymbol{D}_2\boldsymbol{D}_1^{+}, \tag{2.17}$$

where $\boldsymbol{D}_1^{+}$ denotes the pseudo-inverse of the first portion of the dataset. Except for the very special case in which $n_s = n_t$ and the matrix $\boldsymbol{D}_1$ is invertible, the definition of this propagator is not well posed in the classical sense of Hadamard. Moreover, even resorting to the classical Moore–Penrose inverse $\boldsymbol{D}_1^{+} = \boldsymbol{\Psi}_{\mathcal{P}}\boldsymbol{\Sigma}_{\mathcal{P}}^{-1}\boldsymbol{\Phi}_{\mathcal{P}}^{\mathrm{T}}$, the size of this matrix is usually prohibitively large ($\boldsymbol{P} \in \mathbb{R}^{n_s \times n_s}$) and thus the DMD decomposition based on the eigenvalues of $\boldsymbol{P}$ (referred to as exact DMD in Tu *et al.* 2014) is rarely of practical interest. Many variants of the DMD have been developed in the last few years as reviewed in § 1.2. In this work, we only consider the two original DMD algorithms to compare with the proposed mPOD.

The first DMD algorithm, herein denoted as cDMD, reproduces the action of the propagator $\boldsymbol{P}$ with a much smaller matrix $\boldsymbol{\mathcal{C}} \in \mathbb{R}^{(n_t-1)\times(n_t-1)}$ acting on the right side (Rowley *et al.* 2009). This matrix is the companion matrix, defined so that

$$\boldsymbol{\mathcal{C}} = \begin{bmatrix} 0 & 0 & \cdots & 0 & c_1 \\ 1 & 0 & & 0 & c_2 \\ 0 & 1 & & 0 & \vdots \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \cdots & 1 & c_{n_t-1} \end{bmatrix} \Longleftrightarrow \boldsymbol{D}_2 = \boldsymbol{P}\boldsymbol{D}_1 \approx \boldsymbol{D}_1\boldsymbol{\mathcal{C}}. \tag{2.18}$$

The ones in the sub-diagonals shift the columns of $\boldsymbol{D}_1$ to obtain those of $\boldsymbol{D}_2$; the last column contains the set of coefficients $\boldsymbol{c} = [c_1, c_2, \ldots, c_{n_t-1}]^{\mathrm{T}}$ that approximates the last temporal realization $\boldsymbol{d}_{n_t}$ as a linear combination of the previous ones. As an eigenvalue λ of $\boldsymbol{\mathcal{C}}$, with eigenvector $\boldsymbol{v}$ such that $\boldsymbol{\mathcal{C}}\boldsymbol{v} = \lambda\boldsymbol{v}$, is also an eigenvalue of $\boldsymbol{P}$ with eigenvector $\boldsymbol{D}\boldsymbol{v}$, this formulation of the DMD relies on the eigendecomposition of $\boldsymbol{\mathcal{C}}$ and thus the calculation of coefficients $\boldsymbol{c}$. This is a least square minimization problem, extremely sensitive to noise and the choice of the final dataset. Using a $\boldsymbol{QR}$ factorization $\boldsymbol{D}_1 = \boldsymbol{QR}$, the solution can be written as $\boldsymbol{c} = \boldsymbol{R}^{-1}\boldsymbol{Q}^{\mathrm{T}}\boldsymbol{d}_{n_t}$. More robust

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

variants of this method are the optimized DMD by Chen *et al.* (2012) or the sparsity promoting DMD by Jovanović *et al.* (2014) in which the L_2 minimization problem is extended to the entire dataset and not just on the last realization.

The second DMD algorithm, herein denoted as sDMD, avoids the calculation of $\mathcal{C}$ and projects the problem onto a space of much lower dimension (Schmid 2010), in which hopefully only a few eigenvectors are relevant. This is the space spanned by the first r_c dominant POD modes of the dataset.

This propagator is therefore defined as $\boldsymbol{\mathcal{S}} = \tilde{\boldsymbol{\Phi}}_{\mathcal{P}}^{\mathrm{T}} \boldsymbol{P} \tilde{\boldsymbol{\Phi}}_{\mathcal{P}} \in \mathbb{R}^{r_c \times r_c}$, where the tilde denotes an approximation $\tilde{\boldsymbol{\Phi}}_{\mathcal{P}} = [\boldsymbol{\phi}_{\mathcal{P}1}, \ldots, \boldsymbol{\phi}_{\mathcal{P}r_c}]$. Writing the full size propagator using the Moore–Penrose inverse, this projected propagator reads

$$\boldsymbol{\mathcal{S}} = \tilde{\boldsymbol{\Phi}}_{\mathcal{P}}^{\mathrm{T}} \boldsymbol{P} \tilde{\boldsymbol{\Phi}}_{\mathcal{P}} = \tilde{\boldsymbol{\Phi}}_{\mathcal{P}}^{\mathrm{T}} \boldsymbol{D}_2 \tilde{\boldsymbol{\Psi}}_{\mathcal{P}} \tilde{\boldsymbol{\Sigma}}_{\mathcal{P}}^{-1}. \tag{2.19}$$

Observe that (2.19) is a similarity transform, and thus $\boldsymbol{\mathcal{S}}$ and $\boldsymbol{P}$ share the same eigenvalues, only if $\tilde{\boldsymbol{\Phi}}_{\mathcal{P}}$ is a square $n_s \times n_s$ matrix (that is the full POD basis plus its $n_s - n_t$ orthogonal complements are taken into account so that $\tilde{\boldsymbol{\Phi}}_{\mathcal{P}} \tilde{\boldsymbol{\Phi}}_{\mathcal{P}}^{\mathrm{T}} = I$). That is never the case, and one heuristically hopes that at least the first r_c eigenvalues of $\boldsymbol{P}$ are available in $\boldsymbol{\mathcal{S}}$, which implies that $\tilde{\boldsymbol{\Phi}}_{\mathcal{P}} \tilde{\boldsymbol{\Phi}}_{\mathcal{P}}^{\mathrm{T}} \approx I$.

Although numerically more robust, the ill posedness of this POD-based projection is easily illustrated with the test cases in §§ 4 and 5. Specifically, major problems occur when the POD basis captures the energy content of the dataset within a few modes, each of which possibly having several essential frequencies. In this condition, the projected propagator in (2.19) acts on a space which is too small to allow for a sufficient number of frequencies, and the DMD decomposition diverges exponentially fast.

Finally, concerning the mean-shifted form in (2.7), one should notice that the DMD can provide the constant vector $\boldsymbol{\psi}_\mu = \underline{1}/\sqrt{n_t}$ as basis element only if $\lambda = 1$ is an eigenvalue of the propagator ($\boldsymbol{\mathcal{C}}$ for the cDMD, $\boldsymbol{\mathcal{S}}$ for the sDMD). While there is no simple relation to know when this is the case for $\boldsymbol{\mathcal{S}}$, one can see from the characteristic polynomial of $\boldsymbol{\mathcal{C}}$ that this condition sets $\langle \underline{1}, \boldsymbol{c} \rangle = 0$, that is the sum of all the coefficients should be null. This condition is not generally imposed in the calculation of the DMD, so the corresponding basis might not necessarily have a mode to represent the temporal average of data. On the other hand, removing the mean from the dataset results in $\boldsymbol{c} = -\underline{1}$ regardless of the dataset, since the last realization must cancel the summation. In this case, as shown by Chen *et al.* (2012), the solution of the characteristic polynomials are the $n_t - 1$ roots of unity, and the DMD reduces to the DFT.

3. The multi-scale proper orthogonal decomposition (mPOD)

The temporal basis $\boldsymbol{\Psi}_{\mathcal{M}}$ for the multi-scale proper orthogonal decomposition (mPOD) proposed in this work is computed by setting frequency constraints to the classical POD.

The algorithm consists of three major steps. First, the temporal correlation matrix $\boldsymbol{K} = \boldsymbol{D}^{\dagger} \boldsymbol{D}$ is split into the contributions of different scales via multi-resolution analysis (MRA). This can be performed via two-dimensional (2-D) wavelet transform or more generally via an appropriate filter bank that preserves the symmetry of the correlation matrix. Second, each of these contributions is diagonalized as in the classical POD, to obtain the orthonormal basis corresponding to each scale. Third, the bases of each scale are merged into the final temporal basis $\boldsymbol{\Psi}_{\mathcal{M}}$, and the decomposition is completed via (2.6).

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

The theoretical foundation of the algorithm is described in §§ 3.1–3.3. In particular, § 3.1 investigates the link between the frequency content in $\boldsymbol{D}$, in $\boldsymbol{K}$ and in its eigenvectors $\boldsymbol{\Psi}_{\mathcal{P}}$. Section 3.2 introduces the generalized MRA of $\boldsymbol{K}$ via filter banks, discussing its impact on the eigenvectors of different scales and the necessary conditions to keep these mutually orthogonal. Section 3.3 discusses the link between POD, mPOD and DFT, while § 3.4 presents the proposed mPOD algorithm.

3.1. *Mode spectra and correlation spectra*

The multi-resolution analysis (MRA) is introduced in § 3.2 in the frequency domain, that is along the 2-D Fourier transform of the correlation matrix. It is thus interesting to first analyse the standard POD in the frequency domain. At the scope, we denote as $\widehat{\boldsymbol{A}}$ the discrete Fourier transform (DFT) of a matrix $\boldsymbol{A}$, whether this is done along its rows, columns or both. Using (2.13), it is easy to see that the DFT of a column vector is obtained via left multiplication by $\overline{\boldsymbol{\Psi}}_{\mathcal{F}}$, while the DFT of a row vector is a right multiplication by $\overline{\boldsymbol{\Psi}}_{\mathcal{F}}$.

The three Fourier pairs to be linked in this section are related to the time evolution of the data ($\boldsymbol{D}$, row-wise), the temporal structures of the POD modes ($\boldsymbol{\Psi}_{\mathcal{P}}$, column-wise) and the temporal correlation matrix ($\boldsymbol{K}$, over both columns and rows). These are:

$$\widehat{\boldsymbol{D}} = \boldsymbol{D}\overline{\boldsymbol{\Psi}}_{\mathcal{F}} \Longleftrightarrow \boldsymbol{D} = \widehat{\boldsymbol{D}}\boldsymbol{\Psi}_{\mathcal{F}} \tag{3.1a}$$

$$\widehat{\boldsymbol{\Psi}}_{\mathcal{P}} = \overline{\boldsymbol{\Psi}}_{\mathcal{F}}\boldsymbol{\Psi}_{\mathcal{P}} \Longleftrightarrow \boldsymbol{\Psi}_{\mathcal{P}} = \boldsymbol{\Psi}_{\mathcal{F}}\widehat{\boldsymbol{\Psi}}_{\mathcal{P}} \tag{3.1b}$$

$$\widehat{\boldsymbol{K}} = \overline{\boldsymbol{\Psi}}_{\mathcal{F}}\boldsymbol{K}\overline{\boldsymbol{\Psi}}_{\mathcal{F}} \Longleftrightarrow \boldsymbol{K} = \boldsymbol{\Psi}_{\mathcal{F}}\widehat{\boldsymbol{K}}\boldsymbol{\Psi}_{\mathcal{F}}. \tag{3.1c}$$

A first link between the DFT of the dataset (3.1*a*) and that of the POD modes (3.1*b*) is identified by the cross-spectral density matrix, defined and linked to the temporal correlation matrix as follows:

$$\boldsymbol{K}_{\mathcal{F}} = \widehat{\boldsymbol{D}}^{\dagger}\widehat{\boldsymbol{D}} = \boldsymbol{\Psi}_{\mathcal{F}}[\boldsymbol{D}^{\dagger}\boldsymbol{D}]\overline{\boldsymbol{\Psi}}_{\mathcal{F}} = \boldsymbol{\Psi}_{\mathcal{F}}\boldsymbol{K}\overline{\boldsymbol{\Psi}}_{\mathcal{F}} \Longleftrightarrow \boldsymbol{K} = \overline{\boldsymbol{\Psi}}_{\mathcal{F}}\boldsymbol{K}_{\mathcal{F}}\boldsymbol{\Psi}_{\mathcal{F}}. \tag{3.2}$$

This matrix is obtained from the temporal correlation matrix $\boldsymbol{K}$ via a similarity transform (since $\boldsymbol{\Psi}_{\mathcal{F}}\overline{\boldsymbol{\Psi}}_{\mathcal{F}} = \boldsymbol{I}$), and therefore share the same eigenvalues.

To link its eigenvectors to the POD temporal structures, it is important to observe that all the transformed quantities in (3.1*a*)–(3.1*c*) (that is $\boldsymbol{D}$, $\boldsymbol{\Psi}_{\mathcal{P}}$ and $\boldsymbol{K}$) are real, and thus their Fourier transforms are Hermitian symmetric. For a continuous signal $x(t)$, having Fourier transform $\mathcal{F}\{x(t)\} = X(\omega)$, the Hermitian symmetry sets $X(\omega) = \overline{X(-\omega)}$. For a discrete signal, this property can be introduced via a matrix permutation $\boldsymbol{P}_{\pi}$, which flips the spectra along the zero frequency (first column of $\boldsymbol{\Psi}_{\mathcal{F}}$ in (2.13)).

This permutation matrix can be obtained by applying the DFT operator twice:

$$\boldsymbol{P}_{\pi} = \boldsymbol{\Psi}_{\mathcal{F}}\boldsymbol{\Psi}_{\mathcal{F}} = \overline{\boldsymbol{\Psi}}_{\mathcal{F}}\overline{\boldsymbol{\Psi}}_{\mathcal{F}} = \begin{bmatrix} 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & & 0 & 1 \\ 0 & 0 & & 1 & \vdots \\ \vdots & \vdots & \iddots & & \vdots \\ 0 & 1 & 0 & \cdots & 0 \end{bmatrix}. \tag{3.3}$$

Therefore, for any discrete real vector $\boldsymbol{a}$ with discrete Fourier transform $\hat{\boldsymbol{a}} = \boldsymbol{\Psi}_{\mathcal{F}}\boldsymbol{a}$, the Hermitian symmetry can be written as $\hat{\boldsymbol{a}} = \overline{\boldsymbol{P}_{\pi}\hat{\boldsymbol{a}}}$. The fact that the permutation

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

matrix can be obtained by transforming the discrete vector twice reflects the symmetry (or duality) property of the Fourier transform.

Using the Hermitian symmetry on the first two Fourier pairs in (3.1*a*)–(3.1*b*) gives:

$$\widehat{\boldsymbol{\Psi}}_{\mathcal{P}} = \overline{\boldsymbol{P}_{\pi}\widehat{\boldsymbol{\Psi}}_{\mathcal{P}}} \Longleftrightarrow \overline{\widehat{\boldsymbol{\Psi}}_{\mathcal{P}}} = \boldsymbol{P}_{\pi}\widehat{\boldsymbol{\Psi}}_{\mathcal{P}} \Longleftrightarrow \boldsymbol{\Psi}_{\mathcal{F}}\boldsymbol{\Psi}_{\mathcal{P}} = \boldsymbol{P}_{\pi}\widehat{\boldsymbol{\Psi}}_{\mathcal{P}} \tag{3.4a}$$

$$\widehat{\boldsymbol{D}} = \overline{\widehat{\boldsymbol{D}}\boldsymbol{P}_{\pi}} \Longleftrightarrow \overline{\widehat{\boldsymbol{D}}} = \widehat{\boldsymbol{D}}\boldsymbol{P}_{\pi} \Longleftrightarrow \boldsymbol{D}\boldsymbol{\Psi}_{\mathcal{F}} = \widehat{\boldsymbol{D}}\boldsymbol{P}_{\pi}. \tag{3.4b}$$

Introducing the eigenvalue decomposition of $\boldsymbol{K}$ in (3.2) and using (3.4*a*) yields

$$\boldsymbol{K}_{\mathcal{F}} = \boldsymbol{\Psi}_{\mathcal{F}}[\boldsymbol{\Psi}_{\mathcal{P}}\boldsymbol{\Sigma}_{\mathcal{P}}^2\boldsymbol{\Psi}_{\mathcal{P}}^{\dagger}]\overline{\boldsymbol{\Psi}}_{\mathcal{F}} = (\boldsymbol{\Psi}_{\mathcal{F}}\boldsymbol{\Psi}_{\mathcal{P}})\boldsymbol{\Sigma}_{\mathcal{P}}^2(\boldsymbol{\Psi}_{\mathcal{F}}\boldsymbol{\Psi}_{\mathcal{P}})^{\dagger} = (\boldsymbol{P}_{\pi}\widehat{\boldsymbol{\Psi}}_{\mathcal{P}})\boldsymbol{\Sigma}_{\mathcal{P}}^2(\boldsymbol{P}_{\pi}\widehat{\boldsymbol{\Psi}}_{\mathcal{P}})^{\dagger}. \tag{3.5}$$

This shows that the eigenvectors of $\boldsymbol{K}_{\mathcal{F}}$ are the permuted spectra of the POD modes. While the similarity argument is essential in any spectral formulation of the POD (Glauser *et al.* 1987; Citriniti & George 2000; Gordeyev & Thomas 2000; Gudmundsson & Colonius 2011; Sieber *et al.* 2017; Towne *et al.* 2018), (3.5) shows that care must be taken with the conjugation (or the flipping of) the eigenvectors before the inverse DFT, if one seeks to compute POD modes from the cross-spectral density matrix in (3.2).

The DFT of the correlation matrix $\boldsymbol{K}$ in (3.1*c*), on the other hand, is not obtained via similarity transform and is not self-adjoint ($\widehat{\boldsymbol{K}} \neq \widehat{\boldsymbol{K}}^{\dagger}$). Hence, since it does not share the same eigenvalues of $\boldsymbol{K}$ and $\boldsymbol{K}_{\mathcal{F}}$, its diagonalization is of no interest for the purpose of this work. However, because of its relevance on the filtering process, it is important to highlight – by introducing (3.2) into (3.1*c*) – its link to $\boldsymbol{K}_{\mathcal{F}}$:

$$\widehat{\boldsymbol{K}} = \overline{\boldsymbol{\Psi}}_{\mathcal{F}}[\overline{\boldsymbol{\Psi}}_{\mathcal{F}}K_{\mathcal{F}}\boldsymbol{\Psi}_{\mathcal{F}}]\overline{\boldsymbol{\Psi}}_{\mathcal{F}} = \overline{\boldsymbol{\Psi}}_{\mathcal{F}}\overline{\boldsymbol{\Psi}}_{\mathcal{F}}K_{\mathcal{F}} = \boldsymbol{P}_{\pi}\boldsymbol{K}_{\mathcal{F}}. \tag{3.6}$$

3.2. *Decomposing data and decomposing correlations*

It is now of interest to understand how temporal filtering of the dataset influences the frequency content of the POD modes. In particular, we are interested not in just one filter, but an array of filters, constructed to isolate portions of the frequency spectra.

To illustrate the idea, let us consider the Fourier transform $\widehat{\boldsymbol{d}}_p = \boldsymbol{d}_p\overline{\boldsymbol{\Psi}}_{\mathcal{F}}$ of the temporal evolution (row-wise) of the data for a given point spatial location $\boldsymbol{i}_p$, that is $\boldsymbol{d}_p[k] = \boldsymbol{D}[\boldsymbol{i}_p, k]$. We seek to split the spectra $\widehat{\boldsymbol{d}}_p$ into M scales, each retaining a portion with negligible overlapping as pictorially illustrated in figure 1. The frequency bandwidths of these scales $\Delta f_m = f_{m+1}^c - f_m^c$, with $m \in [1, \ldots, M-1]$, are defined by a frequency splitting vector $\boldsymbol{F}_V = [f_1^c, f_2^c, \ldots, f_{M-1}^c]$. This vector is bounded by the minimum ($f_0^c = 0$) and the maximum ($f_M^c = f_s/2$) possible frequencies.

This splitting requires the definition of a filter bank (Vaidyanathanm 1992; Strang & Nguyen 1996; Shukla & Tiwari 2013), consisting of a low-pass filter with cutoff f_1^c, a high-pass filter with cutoff f_{M-1}^c and $M-2$ band-pass filters between these. Following the lossless formulation from multi-resolution analysis (MRA), these filters are constructed from the set of m low-pass filters with transfer function $\boldsymbol{\mathcal{H}}_{\mathcal{L}_m}$. Of these low-pass filters, only the first one ($\boldsymbol{\mathcal{H}}_{\mathcal{L}_1}$) is retained, while the others (shown in figure 1 with dashed line) are only used to build the band-pass filters as complementary differences $\boldsymbol{H}_{\mathcal{H}m} = \boldsymbol{H}_{\mathcal{L}_{m+1}} - \boldsymbol{H}_{\mathcal{L}_m}$, the finest one being $\boldsymbol{H}_{\mathcal{H}_M} = 1 - \boldsymbol{\mathcal{H}}_{\mathcal{L}_M}$. The resulting set of transfer functions (with modulus sketched in figure 1 with a continuous line) is by construction such that $\boldsymbol{H}_{\mathcal{L}_1} + \boldsymbol{H}_{\mathcal{H}_1} + \cdots + \boldsymbol{H}_{\mathcal{H}_{M-1}} = \underline{1}$ and thus

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

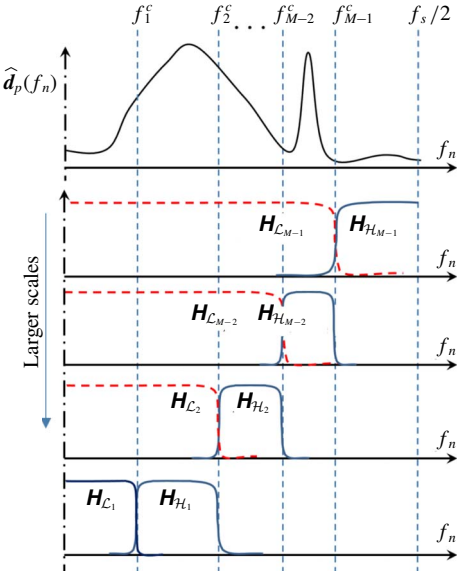


FIGURE 1. (Colour online) Multi-resolution analysis (MRA) of a signal $\boldsymbol{d}_p[k]$ along its spectrum $\widehat{\boldsymbol{d}}_p[n]$. Using a filter bank, the MRA splits the spectrum into a set of M scales. This set consists of a large-scale (transfer function $\boldsymbol{H}_{\mathcal{L}_1}$), a fine-scale (transfer function $\boldsymbol{H}_{\mathcal{H}_{M-1}}$) and $M-2$ intermediate scales (transfer functions from $\boldsymbol{H}_{\mathcal{H}_1}$ to $\boldsymbol{H}_{\mathcal{H}_{M-2}}$). Only the positive portion of the spectrum is shown, the negative one being mirrored along the ordinates.

the entire spectrum of the dataset is retained. The design of the low-pass filters in the filter bank used in this work is presented in appendix A; in what follows, we focus on the algebra of the filtering process and the theoretical foundation of the proposed mPOD.

Considering the filtering in the frequency domain, this operation is performed via simple multiplication (entry by entry), which in matrix notation corresponds to the Hadamard product $\odot$. Moreover, since each of these filters is applied in the entire spatial domain (along each row of $\boldsymbol{D}$), one should copy the transfer functions $\boldsymbol{H}_{\mathcal{L}_m}$ or $\boldsymbol{H}_{\mathcal{H}_m}$ (row vectors $\in \mathbb{C}^{1\times n_t}$) to obtain $n_s \times n_t$ matrices having all the rows equal to $\boldsymbol{H}_{\mathcal{L}_m}$ or $\boldsymbol{H}_{\mathcal{H}_m}$. Let $\boldsymbol{H}'_m$ be the general transfer function (whether low pass or band pass), we denote with an apex the bi-dimensional extension of the 1-D transfer function. The contribution of a scale m from figure 1 in each of the spatial points (the filtered portion of the dataset) is

$$\boldsymbol{D}_m = \overbrace{[\underbrace{(\boldsymbol{D}\overline{\boldsymbol{\Psi}}_{\mathcal{F}})}_{\widehat{\boldsymbol{D}}} \odot \boldsymbol{H}'_m]}^{\widehat{\boldsymbol{D}}_m} \boldsymbol{\Psi}_{\mathcal{F}}. \tag{3.7}$$

The first multiplication is the DFT along the time domain ($\widehat{\boldsymbol{D}}$), the result of the Hadamard multiplication in the square bracket is the filtered spectrum of the data ($\widehat{\boldsymbol{D}}_m$, the frequency contribution of the given scale) and the last right multiplication is the inverse DFT from (3.1*a*). Using (3.7), the correlation matrix $\boldsymbol{K}_m = \boldsymbol{D}_m^{\dagger}\boldsymbol{D}_m$ from

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

each scale contribution reads

$$\boldsymbol{K}_m = \overline{\boldsymbol{\Psi}}_{\mathcal{F}}[(\widehat{\boldsymbol{D}} \odot \boldsymbol{H}'_m)^{\dagger}(\widehat{\boldsymbol{D}} \odot \boldsymbol{H}'_m)]\boldsymbol{\Psi}_{\mathcal{F}} = \overline{\boldsymbol{\Psi}}_{\mathcal{F}}[(\widehat{\boldsymbol{D}}^{\dagger}\widehat{\boldsymbol{D}}) \odot ((\boldsymbol{H}'_m)^{\dagger} \odot \boldsymbol{H}'_m)]\boldsymbol{\Psi}_{\mathcal{F}}, \tag{3.8}$$

where the associative and commutative properties of the Hadamard product are used to move the transfer functions. Introducing the cross-spectral density matrix $\boldsymbol{K}_{\mathcal{F}}$ in (3.2) and writing the 2-D transfer function as $\underline{\boldsymbol{\mathcal{H}}}_m = (\boldsymbol{H}'_m)^{\dagger} \odot \boldsymbol{H}'_m$ yields

$$\boldsymbol{K}_m = \overline{\boldsymbol{\Psi}}_{\mathcal{F}} \underbrace{[\boldsymbol{K}_{\mathcal{F}} \odot \underline{\boldsymbol{\mathcal{H}}}_m]}_{\boldsymbol{K}_{\mathcal{F}m}} \boldsymbol{\Psi}_{\mathcal{F}} = \boldsymbol{\Psi}_{\mathcal{F}} \underbrace{[\widehat{\boldsymbol{K}} \odot \underline{\boldsymbol{\mathcal{H}}}_m]}_{\widehat{\boldsymbol{K}}_m} \boldsymbol{\Psi}_{\mathcal{F}}, \tag{3.9}$$

having introduced (3.6) in the last step.

This important result lays the foundation of the mPOD. First, the expression on the left of (3.9) is a similarity transform: the filtered cross-spectral density $\boldsymbol{K}_{\mathcal{F}m} = \boldsymbol{K}_{\mathcal{F}} \odot \underline{\boldsymbol{\mathcal{H}}}_m$ shares the same eigenvalues of the correlation of filtered data $\boldsymbol{K}_m = \boldsymbol{D}_m^{\dagger}\boldsymbol{D}_m$. Second, introducing the eigenvalue decomposition of $\boldsymbol{K}_m$ as in (3.5), shows that – up to a permutation – the eigenvectors of $\boldsymbol{K}_{\mathcal{F}}$ are the DFT of the POD modes of $\boldsymbol{D}_m$:

$$\boldsymbol{K}_{\mathcal{F}m} = \boldsymbol{\Psi}_{\mathcal{F}} K_m \overline{\boldsymbol{\Psi}}_{\mathcal{F}} = \boldsymbol{\Psi}_{\mathcal{F}}[\boldsymbol{\Psi}_{\mathcal{P}m}\boldsymbol{\Sigma}^2_{\mathcal{P}m}\boldsymbol{\Psi}^{\dagger}_{\mathcal{P}m}]\overline{\boldsymbol{\Psi}}_{\mathcal{F}} = (\boldsymbol{P}_{\pi}\widehat{\boldsymbol{\Psi}}_{\mathcal{P}m})\boldsymbol{\Sigma}^2_{\mathcal{P}m}(\boldsymbol{P}_{\pi}\widehat{\boldsymbol{\Psi}}_{\mathcal{P}m})^{\dagger}, \tag{3.10}$$

having introduced (3.4*a*) in the last step. The impact of the filter on the POD modes is revealed by the diagonal entries of $\boldsymbol{K}_{\mathcal{F}m}$ and $\boldsymbol{K}_{\mathcal{F}}$, which are summations of non-negative real numbers. From (3.10), using (3.4*a*), these are:

$$\boldsymbol{K}_{\mathcal{F}m}[i,j] = \boldsymbol{K}_{\mathcal{F}}[i,j] \odot \underline{\boldsymbol{\mathcal{H}}}_m[i,j] = \sum_{r=1}^{n_t} \sigma^2_{\mathcal{P}mr}\overline{\widehat{\boldsymbol{\psi}}}_{\mathcal{P}mr}[i]\overline{\widehat{\boldsymbol{\psi}}}^{\dagger}_{\mathcal{P}mr}[j]. \tag{3.11}$$

If the filter at a given scale m removes one of the entries along the diagonal of $\boldsymbol{K}_{\mathcal{F}}$ ($\underline{\boldsymbol{\mathcal{H}}}_m[i,i] = 0$), the associated frequency $\widehat{\boldsymbol{\psi}}_{\mathcal{P}mr}[i]$ cannot be present in any of the r POD modes of the filtered data: the POD operating at this scale is constrained within the frequencies allowed by the transfer function.

Moreover, while the expression on the right side of (3.9) is not a similarity transform, it shows that the correlation matrix associated with the filtered data $\boldsymbol{K}_m$ can be computed by a simple 2-D filtering of the original correlation matrix, saving considerable computational time with respect to n_s filters along the rows of $\boldsymbol{D}$. Equation (3.9) shows that the MRA of the dataset can be carried out via 2-D MRA of the temporal correlation matrix.

The four possible structures for the 2-D transfer functions $\underline{\boldsymbol{\mathcal{H}}}_m$ that arise from a combination of the 1-D kernels in the time domain are illustrated graphically in figure 2. For a given scale, composed of a low frequency and a high frequency $\boldsymbol{D}_m = \boldsymbol{D}_{\mathcal{L}_m} + \boldsymbol{D}_{\mathcal{H}_m}$, these correspond to the contributions of four temporal correlation matrices:

$$\begin{aligned}\boldsymbol{K}_m = \boldsymbol{D}_m^{\dagger}\boldsymbol{D}_m &= \boldsymbol{D}^{\dagger}_{\mathcal{L}_m}\boldsymbol{D}_{\mathcal{L}_m} + \boldsymbol{D}^{\dagger}_{\mathcal{H}_m}\boldsymbol{D}_{\mathcal{H}_m} + \boldsymbol{D}^{\dagger}_{\mathcal{L}_m}\boldsymbol{D}_{\mathcal{H}_m} + \boldsymbol{D}^{\dagger}_{\mathcal{H}_m}\boldsymbol{D}_{\mathcal{L}_m} \\ &= \boldsymbol{\Psi}_{\mathcal{F}}[\widehat{\boldsymbol{K}} \odot \underline{\boldsymbol{\mathcal{H}}}_{\mathcal{L}m} + \widehat{\boldsymbol{K}} \odot \underline{\boldsymbol{\mathcal{H}}}_{\mathcal{H}m} + \widehat{\boldsymbol{K}} \odot \underline{\boldsymbol{\mathcal{H}}}_{\mathcal{L}\mathcal{H}m} + \widehat{\boldsymbol{K}} \odot \underline{\boldsymbol{\mathcal{H}}}_{\mathcal{H}\mathcal{L}m}]\boldsymbol{\Psi}_{\mathcal{F}}.\end{aligned} \tag{3.12}$$

The 'pure' terms $\boldsymbol{D}^{\dagger}_{\mathcal{L}_m}\boldsymbol{D}_{\mathcal{L}_m}$ and $\boldsymbol{D}^{\dagger}_{\mathcal{H}_m}\boldsymbol{D}_{\mathcal{H}_m}$ can be obtained by first filtering the data (3.7) and then computing the related correlations; the mixed terms can only be revealed by a MRA of the full correlation matrix $\boldsymbol{K}_m$.

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

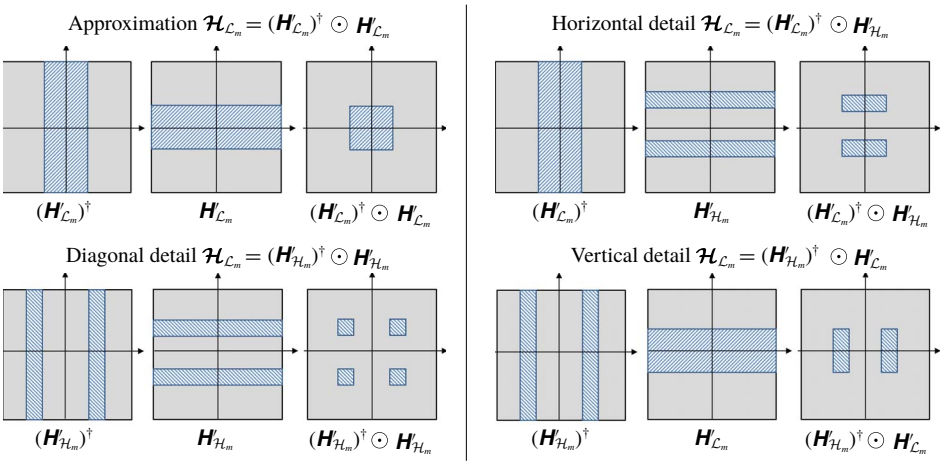


FIGURE 2. (Colour online) The magnitude of the four possible 2-D transfer functions $\underline{\mathcal{H}}_m$ that can be obtained by combining the low-pass ($\boldsymbol{H}_{\mathcal{L}_m}$) and the high-pass ($\boldsymbol{H}_{\mathcal{H}_m}$) filters associated with a given m scale of $\boldsymbol{K}$, following the multi-scale architecture in figure 1. The areas where the transfer function is ~ 1 and ~ 0 are indicated with blue dashed lines and full grey respectively.

It is worth noticing that the filter bank architecture described in this section is a generalized version of the MRA via the discrete (dyadic) wavelet transform (DWT) proposed in previous works (Mendez *et al.* 2018*b*; Mendez, Balabane & Buchlin 2018*a*). In the 2-D wavelet terminology, the four terms in (3.12), with spectral band-pass region illustrated in figure 2, correspond to the approximation and the diagonal, horizontal and vertical details of each scale.

Particular emphasis should be given to the 'pure' terms $\boldsymbol{D}^\dagger_{\mathcal{L}_m}\boldsymbol{D}_{\mathcal{L}_m}$ and $\boldsymbol{D}^\dagger_{\mathcal{H}_m}\boldsymbol{D}_{\mathcal{H}_m}$. These contributions have no frequency overlapping and their eigenspaces are orthogonal complements by construction: their set of eigenvectors can be used to assemble a complete and orthogonal basis as illustrated in § 3.4. The mixed terms $\boldsymbol{D}^\dagger_{\mathcal{L}_m}\boldsymbol{D}_{\mathcal{H}_m}$ and $\boldsymbol{D}^\dagger_{\mathcal{H}_m}\boldsymbol{D}_{\mathcal{L}_m}$, on the other hand, generates eigenspaces that can potentially overlap with those of the 'pure' contributions and must be disregarded in the MRA of the correlation matrix.

3.3. *The mPOD as bridge between POD and DFT*

Neglecting the mixed terms prevents the full reconstruction of the correlation matrix, but not the full reconstruction of the dataset since the final temporal basis is complete. The consequences of this operation can be analysed while studying how the choice of the frequency splitting vector $\boldsymbol{F}_V$ allows the mPOD to move from the energy optimality of the POD to the spectral purity of the DFT.

Figure 3 shows the band-pass region from the 2-D spectra of the correlation matrix $\widehat{\boldsymbol{K}}$ for four choices of $\boldsymbol{F}_V$. The case (a) corresponds to a single scale mPOD (empty $\boldsymbol{F}_V$), that is the standard POD. All the eigenvectors $\boldsymbol{\psi}_{\mathcal{P}_r}$ are allowed to span the entire set of discrete frequencies available. This configuration corresponds to perfect reconstruction of the correlation matrix and maximum time localization capabilities: if an impulse $\delta(t_k)$ is produced at a certain time, the POD is free to choose an impulse (which has a frequency spectra spanning the entire frequency axis)

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

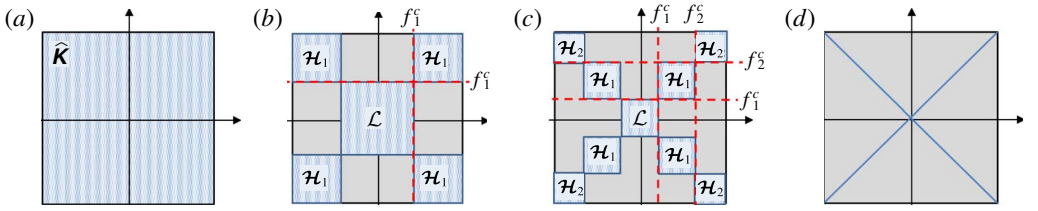


FIGURE 3. (Colour online) Partition of the correlation matrix $\widehat{\boldsymbol{K}}$ for different choices of the splitting vector $\boldsymbol{F}_V$ in the MRA step of the mPOD. The same colour notation of figure 2 is used for the band-pass/band-stop regions. Case (*a*) consists of a single scale, thus the mPOD becomes a POD. Cases (*b*,*c*) consist of two and three scales, respectively. Case (*d*) shows the limit of $n_t/2$ scales, in which the mPOD becomes a DFT.

as one of its temporal structures. On the other hand, all the POD modes share the full set of frequencies, potentially leading to spectral mixing, and phenomena spanning the entire frequency and energy spectra (such as random noise) are spread over all the modes.

The case (*b*) corresponds to a two-scale mPOD, that is with $\boldsymbol{F}_V = f_1^c$. Each scale has its own correlation matrix ($\boldsymbol{K}_{\mathcal{L}}$ and $\boldsymbol{K}_{\mathcal{H}_1}$) with no common frequencies. Phenomena occurring at these two scales are spectrally separated and necessarily assigned to different modes. Moreover, the amount of random noise in each mode is distributed proportionally to the spectral bandwidth of the corresponding scale and becomes less important. Part of the information on the correlation matrix is lost, but the eigenspaces of the two scales, together, span the entire $\mathbb{R}^{n_t}$ space and thus preserve the decomposition convergence. On the other hand, the temporal localization capabilities are reduced: as no eigenvector is allowed to span the full frequency range, no temporal structure can capture impulses nor other phenomena requiring the entire frequency range (e.g. shocks).

As the number of scales is increased, their spectral bandwidth reduces and the decomposition becomes spectrally more localized, at the cost of an increased loss of information in the correlation and a reduction of time localization capabilities. The case (*c*) has three scales, with $\boldsymbol{F}_V = [f_1^c, f_2^c]$, while the case (*d*) corresponds to the limit $\boldsymbol{F}_V \to \{rf_s/(\Delta t n_t)\}_{r=1}^{n_t-1}$, with $r \in [1, \ldots, n_t - 1]$. In this limiting case, every scale is constrained to a single frequency and the mPOD reduces to the DFT, inheriting its finest spectral resolution and its limits (windowing problems, Gibbs phenomena, no time localization and poor convergence, Harris 1978, Daubechies 1992, Folland 2009).

3.4. *The mPOD algorithm*

The steps of the proposed mPOD algorithm are listed in the Algorithm 1. As for any data-driven decomposition, the first step consists in assembling the data matrix as in (2.1). The second step consists in computing the temporal correlation matrix $\boldsymbol{K}$ and its Fourier transform $\widehat{\boldsymbol{K}}$ from (3.1*c*). From the analysis of the frequency content in the correlation spectra $\widehat{\boldsymbol{K}}$, the third step consists in defining the frequency splitting vector $\boldsymbol{F}_V$ and in constructing the set of associated transfer functions.

This choice is the fundamental step which will set the spectral constraints to the mPOD modes. The frequency splitting could be user-defined if prior knowledge on the investigated data is available, or automatically computed by centring the band-pass windows of each scale on the dominant peaks in the diagonal of $\widehat{\boldsymbol{K}}$ and splitting

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

Algorithm 1 Multi-scale proper orthogonal decomposition of a dataset $\boldsymbol{D} \in \mathbb{R}^{n_s \times n_t}$.

Input: set of n_t snapshots of a dataset $\boldsymbol{D}(\boldsymbol{x}_i, t_k)$ on a Cartesian grid $\boldsymbol{x}_i \in \mathbb{R}^{n_x \times n_y}$

1: *Reshape snapshots into column vectors* $\boldsymbol{d}_k[\boldsymbol{i}] \in \mathbb{R}^{n_s \times 1}$ *and assemble* $\boldsymbol{D}[i, k]$ *in (2.1)*
2: *Compute time correlation matrix* $\boldsymbol{K} = \boldsymbol{D}^{\dagger}\boldsymbol{D}$ *and* $\widehat{\boldsymbol{K}} = \overline{\boldsymbol{\Psi}}_{\mathcal{F}}\boldsymbol{K}\overline{\boldsymbol{\Psi}}_{\mathcal{F}}$ *in (3.1c)*
3: *Prepare the filter bank* $\boldsymbol{\mathcal{H}}_{\mathcal{L}_1}, \ldots, \boldsymbol{\mathcal{H}}_{\mathcal{H}_M}$*, and split* $\boldsymbol{K}$ *into M contributions via (3.13)*
4: *Diagonalize each 'pure' term* $\boldsymbol{K}_m = \boldsymbol{\Psi}_m \boldsymbol{\Sigma}_m^2 \boldsymbol{\Psi}_m^{\mathrm{T}}$ *as in (3.14)*
5: *Sort the contribution of all the scales into* $\boldsymbol{\Psi}^0_{\mathcal{M}}$ *as in (3.15)*
6: *Enforce orthogonality via* $\boldsymbol{QR}$ *factorization,* $\boldsymbol{\Psi}^0_{\mathcal{M}} = \boldsymbol{\Psi}_{\mathcal{M}}\boldsymbol{R} \rightarrow \boldsymbol{\Psi}_{\mathcal{M}} = \boldsymbol{\Psi}^0_{\mathcal{M}}\boldsymbol{R}^{-1}$
7: *Compute the spatial basis* $\boldsymbol{\Phi}_{\mathcal{M}} = \boldsymbol{D}\boldsymbol{\Psi}_{\mathcal{M}}\boldsymbol{\Sigma}_{\mathcal{M}}^{-1}$ *from (2.6) and sort the results in descending order of energy contribution.*

Output: spatial $\boldsymbol{\Phi}_{\mathcal{M}}$ and temporal $\boldsymbol{\Psi}_{\mathcal{M}}$ structures with corresponding amplitudes $\boldsymbol{\Sigma}_{\mathcal{M}}$ of the n_t mPOD modes.

the band-pass regions accordingly. In this step, only the 'pure' (approximation and diagonal detail) terms of the 2-D transfer functions are constructed, that is $\boldsymbol{\mathcal{H}}_{\mathcal{L}_1} = (\boldsymbol{H}'_{\mathcal{L}_1})^{\dagger} \odot \boldsymbol{H}'_{\mathcal{L}_1}$ and $\boldsymbol{\mathcal{H}}_{\mathcal{H}_m} = (\boldsymbol{H}'_{\mathcal{H}_m})^{\dagger} \odot \boldsymbol{H}'_{\mathcal{H}_m}$. Disregarding the mixed terms (horizontal and vertical details) in each scale, the correlation matrix is approximated as:

$$\boldsymbol{K} \approx \boldsymbol{\Psi}_{\mathcal{F}}[\widehat{\boldsymbol{K}} \odot \boldsymbol{\mathcal{H}}_{\mathcal{L}_1}]\boldsymbol{\Psi}_{\mathcal{F}} + \sum_{m=1}^{M} \boldsymbol{\Psi}_{\mathcal{F}}[\widehat{\boldsymbol{K}} \odot \boldsymbol{\mathcal{H}}_{\mathcal{H}_m}]\boldsymbol{\Psi}_{\mathcal{F}} \approx \boldsymbol{K}_{\mathcal{L}_1} + \sum_{m=1}^{M-1} \boldsymbol{K}_{\mathcal{H}_m}. \tag{3.13}$$

Each of these contributions is a symmetric, real and positive definite matrix, equipped with its orthonormal eigenspace, computed in the fourth step

$$\boldsymbol{K} \approx \boldsymbol{\Psi}_{\mathcal{L}_1} \boldsymbol{\Sigma}^2_{\mathcal{L}_1} \boldsymbol{\Psi}^{\mathrm{T}}_{\mathcal{L}_1} + \sum_{m=1}^{M-1} \boldsymbol{\Psi}_{\mathcal{H}m} \boldsymbol{\Sigma}^2_{\mathcal{H}_m} \boldsymbol{\Psi}^{\mathrm{T}}_{\mathcal{H}_m}. \tag{3.14}$$

If the frequency overlap is identically null, the eigenspaces of these pure terms are orthogonal complements and therefore there are at most n_t non-zero eigenvalues among all the scales. This is a major difference between the mPOD and other multi-scale methods such as continuous wavelet transform (CWT), in which the temporal basis is constructed by shifting and dilating a 'mother' function, or the multi-resolution DMD (mrDMD) proposed by Kutz *et al.* (2016*b*), in which the temporal basis is constructed by performing DMD on different portions of the datasets. These decompositions potentially produce high redundancy and poor convergence since the basis is larger than n_t. Moreover, each of the basis elements in the mPOD exists over the entire time domain (although they could be null in an arbitrarily large portion of it), while CWT or mrDMD produce different bases for different portions of the time domain, leading to decompositions more complicated than those analysed in this work.

Because of the limited size of the filter impulse response (see appendix A) of each filter, the number of non-zero eigenvalues among the scales is in practice usually slightly larger than n_t. Therefore, only the first n_t dominant eigenvectors are selected. This is done in the fifth step, by first sorting the full set of eigenvalues $[\mathrm{diag}(\boldsymbol{\Sigma}_{\mathcal{L}_1}), \mathrm{diag}(\boldsymbol{\Sigma}_{\mathcal{H}_1}), \ldots, \mathrm{diag}(\boldsymbol{\Sigma}_{\mathcal{H}_{M-1}})]$ in decreasing order and storing the required permutation matrix $\boldsymbol{P}_{\boldsymbol{\Sigma}}$. This matrix is then used to re-arrange the eigenvector matrices

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

	POD	DFT	sDMD	cDMD	mPOD ($M=3$)	mPOD ($M=10$)	mPOD ($M=50$)
Test case 1 (§ 4)	1.7	5.1	—	15.1	2.3	2.9	6.3
Test case 2 (§ 5)	10.1	14.5	53.45	63.3	24.1	54.1	235
Test case 3 (§ 6)	15.8	48.8	—	113.0	51.6	135	595

TABLE 1. CPU time (in seconds) for the test cases analysed in this work and the decompositions tested. The analysis is carried out using Matlab R2018b on a Intel(TM) i7-3770 @3.40 GHz processor with a 32 GB RAM.

$\boldsymbol{\Psi}_{\mathcal{L}}, \ldots, \boldsymbol{\Psi}_{\mathcal{H}m}$ into one single matrix (temporal basis) regardless of their scale of origin

$$\boldsymbol{\Psi}^0_{\mathcal{M}} = [\boldsymbol{\Psi}_{\mathcal{L}_1}, \boldsymbol{\Psi}_{\mathcal{H}_1}, \boldsymbol{\Psi}_{\mathcal{H}_2}, \ldots, \boldsymbol{\Psi}_{\mathcal{H}_{M-1}}]\boldsymbol{P}_{\boldsymbol{\Sigma}}. \tag{3.15}$$

As discussed in § 3.2, the temporal matrix (3.15) is orthonormal if a perfect spectral separation is achieved. This is a major difference between the mPOD and the SPOD by Sieber *et al.* (2016), in which the filtering procedure can eventually result in a non-orthogonal temporal basis, or the mrDMD by Kutz *et al.* (2016*b*), in which the loss of orthogonality can be produced by growing/decaying modes.

Due to the finite size of the transition band of the filters, this is generally not the case in practice, and a minor loss of orthogonality usually occurs. The sixth step treats this imperfection, polishing this temporal structure via a reduced $\boldsymbol{QR}$ factorization $\boldsymbol{\Psi}^0_{\mathcal{M}} = \boldsymbol{\Psi}_{\mathcal{M}}\boldsymbol{R}$, so that $\boldsymbol{\Psi}_{\mathcal{M}} = \boldsymbol{\Psi}^0_{\mathcal{M}}\boldsymbol{R}^{-1} \in \mathcal{R}^{n_t \times n_t}$. This compensates for the losses of orthogonality and the upper triangular matrix $\boldsymbol{R}$ offers an indication of the quality of the filtering process, being close to the identity when the spectral overlapping is negligible. Finally, in the last step, the decomposition is completed via the projection and the normalization in (2.6), as for any orthogonal decomposition.

To conclude the presentation of the mPOD algorithm, table 1 collects the CPU time for all the decompositions used in this work and the test cases in §§ 4–6. The POD is naturally the fastest requiring only an eigenvalue decomposition and a projection, with no need for amplitude normalization and sorting. The DFT is slightly more expensive, mostly due to the normalization and the sorting steps, while the DMD requires the additional cost of a minimization problem or an SVD depending on the chosen algorithm. The computational cost of the mPOD is shown for $M = 3, 10, 50$ and drastically increases with the number of scales considered.

While the cost of the MRA is usually negligible when using an FFT-based formulation in the frequency domain, the major additional cost is in the diagonalizations for each scale. For $M = 3$, the CPU time for the proposed mPOD algorithm is well below the CPU time of the DMD and become comparable to it for $M \sim 10$. For a much larger number of scales (see table for $M = 50$), the mPOD requires a more efficient implementation of the diagonalization of the various scales. An improved algorithm which performs this step in the frequency domain is under development and will be presented in future work.

4. Example I: illustrative synthetic test

4.1. *Dataset description*

As a first example, we propose a simple synthetic test case which is already in the decomposed form of (2.5) and yet capable of severely challenging standard decompositions. The spatial domain consists of a square Cartesian grid of $n_s =$

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

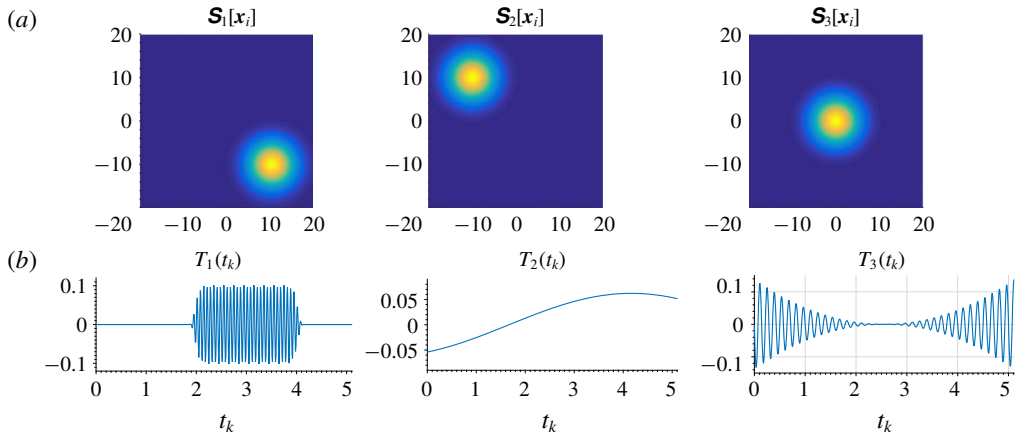


FIGURE 4. (Colour online) Spatial (*a–c*) and corresponding temporal (*d–f*) structures of the modes considered in the simplified test case. The amplitude of each contribution is computed so as to give the energy content. An animation of this test case is available as supplemental Movie 1 available at https://doi.org/10.1017/jfm.2019.212.

256×256 points over a domain $\boldsymbol{x}_i \in [-20, 20] \times [-20, 20]$, while the time domain spans $t_k \in [0, 5.11]$ with a total of $n_t = 256$ points with sampling frequency $f_s = 100$. The dataset is composed of the sum of three modes, having spatial and temporal structures as illustrated in figure 4. These modes consist of identical Gaussians with standard deviations of $\sigma = 5$, located in $\boldsymbol{x}_1 = [10, -10]$, $\boldsymbol{x}_2 = [-10, 10]$ and $\boldsymbol{x}_3 = [0, 0]$, pulsing as

$$T_1(t_k) = A_1 \sin(2\pi f_1 t_k) \exp[0.5(t_k - 3)^{20}] \tag{4.1a}$$

$$T_2(t_k) = A_2 \sin(2\pi f_2 t_k - \pi/3) \tag{4.1b}$$

$$T_3(t_k) = A_3 \sin(2\pi f_3 t_k)(t_k - 2.55)^2. \tag{4.1c}$$

The first mode has a smoothed square box modulation, the second has a harmonic with period longer than the observation time and the third has a parabolic modulation. The amplitudes A_k are chosen so as to equal each contribution, that is $A_r = 1/\|S_r T_r^T\|_2$, having reshaped both spatial and temporal structures as column vectors. The three frequencies are $f_1 = 15$, $f_2 = 0.1$ and $f_3 = 7$. The resulting dataset $\boldsymbol{D}[\boldsymbol{i}, k]$ is a matrix of size $65\,536 \times 256$ of rank $rk(\boldsymbol{D}) = 3$.

4.2. *Modal analysis*

In this test case, we compare the performances of the DFT, POD, cDMD and mPOD. The sDMD algorithm, based on the reduced propagator $\boldsymbol{\mathcal{S}}$ in (2.19), is severely challenged by the rank deficiency of this dataset. Since $rk(\boldsymbol{D}) = 3$, one has $\boldsymbol{\mathcal{S}} \in \mathbb{R}^{3\times3}$, that is a set of eigenvalues incapable of reproducing the data: each of the dominant frequencies f_1, f_2, f_3 requires two complex conjugate modes and thus at least a 6×6 propagator. Low-dimensional propagators $\boldsymbol{\mathcal{S}}$ that are inadequately too small occur for datasets which are close to being rank deficient, in which few coherent patterns occur at different frequencies.

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

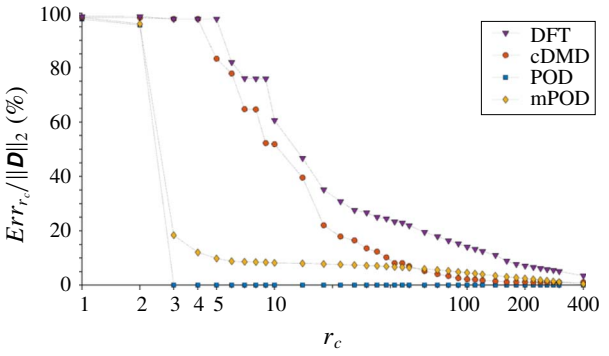


FIGURE 5. (Colour online) Convergence of the POD, cDMD, mPOD and DFT for the illustrative synthetic test case in terms of L^2 relative error, defined as in (4.2).

For all the decomposition, the L^2 convergence is calculated as

$$\frac{Err[r_c]}{\|\boldsymbol{D}\|_2} = \frac{\left\| \boldsymbol{D} - \sum_{r=1}^{r_c} \sigma_r \boldsymbol{\phi}_r \boldsymbol{\psi}_r \right\|_2}{\|\boldsymbol{D}\|_2}, \tag{4.2}$$

where all the modal contributions are first sorted in descending order $\sigma_a > \sigma_b, \forall a > b$. For the POD, (4.2) reduces to $\mathrm{Err}[r_c]/\|\boldsymbol{D}\|_2 = \sigma_{\mathcal{P}r_{c+1}}/\sigma_{\mathcal{P}_1}$ since the L^2 norm of a matrix is equivalent to its largest singular value (see (2.11)). This is not the case for other decompositions, for which the spatial basis is not orthogonal.

The convergence results of this test case are shown in figure 5. As expected, the POD produces no reconstruction error when $r_c \geqslant 3$, while the convergence of frequency based methods is considerably weaker. The finite duration of the first mode, the aperiodicity of the second and the nonlinear modulation of the third are not easily derivable as a linear combination of infinite harmonics, and the convergence of the DFT requires up to 150 modes to bring the relative error below 10 %. The DMD performs better than the DFT, although yet not comparable to the POD. The mPOD of this synthetic test case is performed with $M = 4$, $\boldsymbol{F}_V = [0.4, 10, 20]$ and leads to a convergence much closer to that of the POD than the harmonic decompositions. To analyse the detection capabilities of these decompositions, several representative modes for each are now considered. First, the energy contributions of the DFT ($\mathrm{diag}(\boldsymbol{\Sigma}_{\mathcal{F}})$) and the cDMD ($\mathrm{diag}(\boldsymbol{\Sigma}_{\mathcal{D}})$) are mapped in the complex plane in figure 6. By definition, all the DFT modes lie in the unit circle and are Hermitian symmetric; DMD eigenvalues can lay inside the unit circle (if $|\lambda_r| < 1$) and exponentially decay, or outside (if $|\lambda_r| > 1$) and exponentially grow.

Although the two spectra are similar, this additional degree of freedom gives to the DMD better convergence and reduced windowing problems, especially on the lower frequency. The spatial structures for the first 6 odd DFT modes (the even one are phase-shifted copies) are shown in figure 7. The amplitude of these modes is indicated with a square marker in figure 6 while their temporal structures, being simple harmonics, are not shown. The frequency resolution of the DFT is of the order of $f_s/n_t \approx 0.2$, with the leading modes pulsing at $f_n = [0.19, 7.03, 15.04]$. The spectral leakage is more evident at the lowest frequency, for which a full period of

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

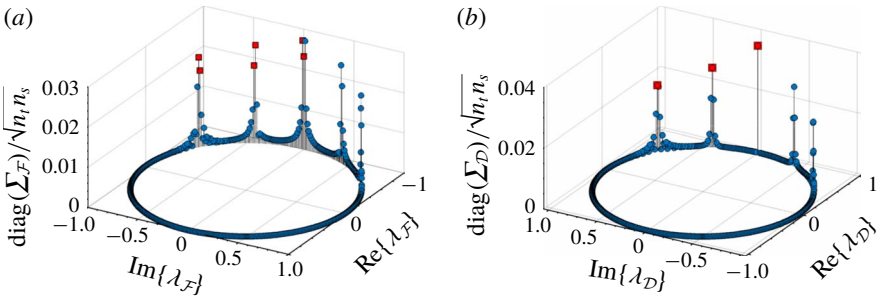


FIGURE 6. (Colour online) Amplitude spectra of the DFT (*a*) and cDMD (*b*) spectra mapped in the complex plane. Both decompositions have temporal evolutions distributed along the unit circle. For the DFT this is true by definition; for the DMD this is the result from the eigenspectra of the companion matrix assembled from the data.

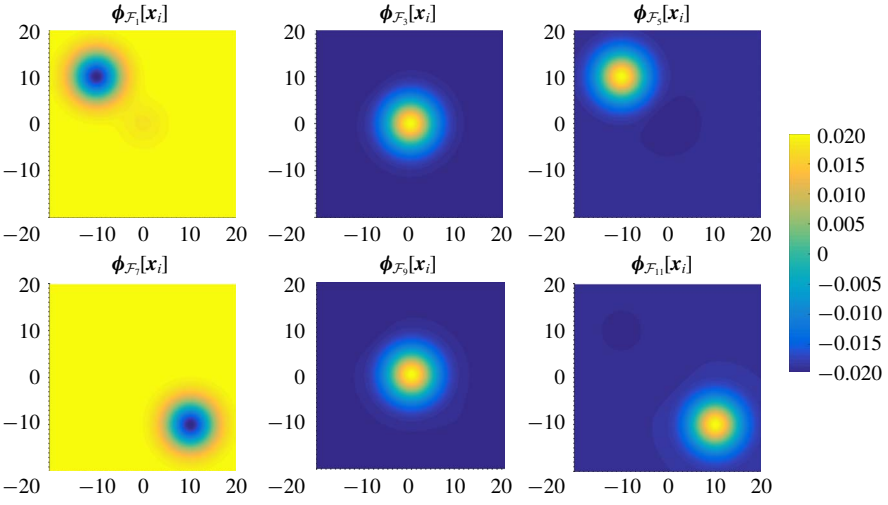


FIGURE 7. (Colour online) Spatial structures $\boldsymbol{\phi}_{\mathcal{F}}[\boldsymbol{x}_i]$ of the first six DFT modes, with amplitude indicated in figure 6 with a square marker.

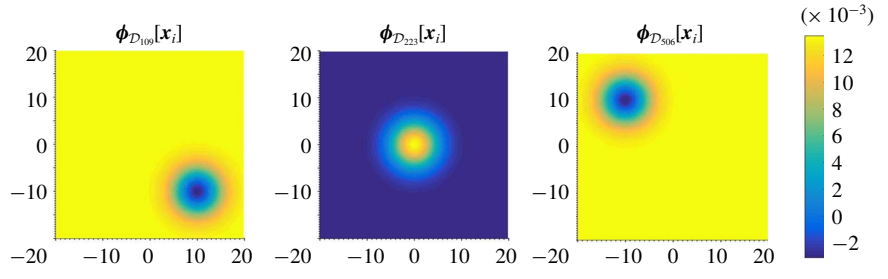


FIGURE 8. (Colour online) Spatial structures $\boldsymbol{\phi}_{\mathcal{D}}[\boldsymbol{x}_i]$ of the first three cDMD modes, with amplitude indicated in figure 6 with a square marker.

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

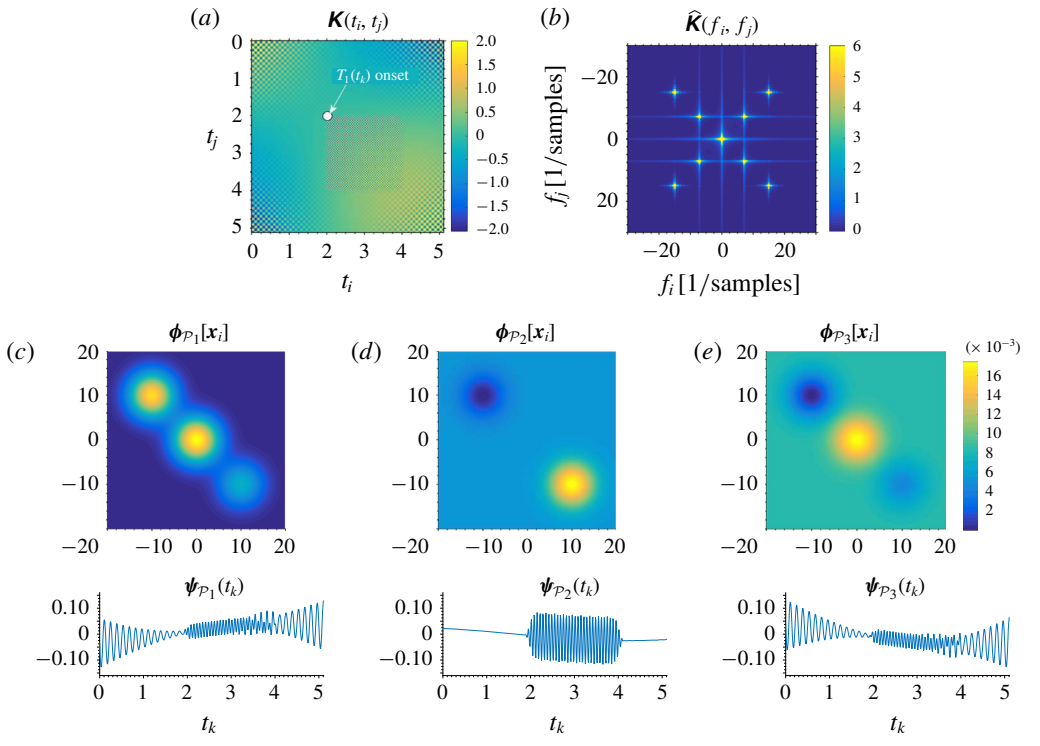


FIGURE 9. (Colour online) Results from the POD of the synthetic test case. (*a*,*b*) Shows the normalized temporal correlation matrix $\boldsymbol{K} = \boldsymbol{D}^{\dagger}\boldsymbol{D}/n_t$ (*a*) and corresponding DFT $\widehat{\boldsymbol{K}} = \boldsymbol{\Psi}_{\mathcal{F}}\boldsymbol{K}\boldsymbol{\Psi}_{\mathcal{F}}/n_t$ (*b*). (*c*–*e*) Shows the three non-zero POD modes, including their spatial structures $\boldsymbol{\phi}_{\mathcal{P}_r}$ and temporal structures $\boldsymbol{\psi}_{\mathcal{P}_r}$.

the oscillation is not captured. These modes have no time localization capabilities: none of these can locate in time the identified structures.

This localization limit, stemming from the Heisenberg uncertainty principle (see Daubechies 1992) is the main limit responsible for the poor convergence, which in turns results in a high redundancy in the spatial modes: many modes (besides the conjugate pairs) have identical spatial structures (e.g. $\boldsymbol{\phi}_{\mathcal{F}_3} \approx \boldsymbol{\phi}_{\mathcal{F}_9}$, $\boldsymbol{\phi}_{\mathcal{F}_5} \approx -\boldsymbol{\phi}_{\mathcal{F}_{11}}$) associated, in the time domain, with harmonic corrections to represent finite duration events.

The DMD modes, having a complex frequency, give a valid alternative when describing events that vanish from the initial condition, but not for events that occur after the initial step. A natural way of overcoming such limit is to break the data into multiple windows and perform the DMD in each of these as in the mrDMD (Kutz *et al.* 2016*b*), at the cost of significantly increasing the decomposition complexity as discussed in § 3.4.

Even in its simplest formulation, the cDMD improves the convergence and has better frequency resolution than the DFT. The cDMD spectrum highlights the three dominant modes pulsing at $f_n = [0.101, 7.00, 15.09]$, and the corresponding spatial structures, shown in figure 8, correctly localize in space these pulsations. The time localization, of course, remains impossible as for any frequency-based method.

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

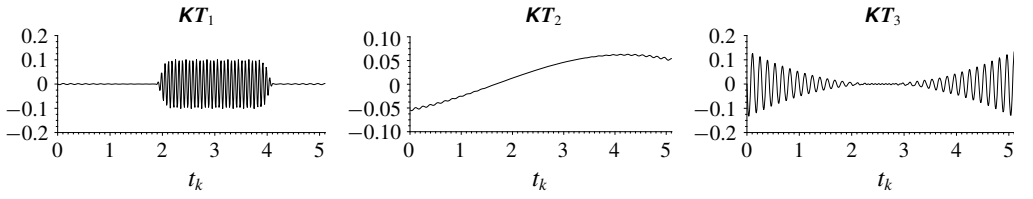


FIGURE 10. Test for POD uniqueness for the proposed test case: the temporal structures in figure 4 are almost eigenvectors of the temporal correlation matrix, since $\boldsymbol{KT}_k \sim \lambda_k \boldsymbol{T}_k$.

By contrast, the POD avoids both the time localization and the convergence problems of DFT and DMD, but produces a severe spectral mixing between different modes. The results of the POD are collected in figure 9, which includes the temporal correlation matrix $\boldsymbol{K}$ and its DFT $\widehat{\boldsymbol{K}}$, as well as the three identified modes. The first carries the information on the time localization of each event, as the superimposition of different patterns is visible (the onset of the temporal evolution $\boldsymbol{T}_1$ is marked in the figure). The second carries the information on the frequency localization, with the three dominant harmonics producing peaks mirrored along the centre in both the diagonal and the anti-diagonal. The spectral leakage produces cross-like patterns centred in these peaks, with larger extension at the lowest frequencies (affected by the strongest windowing effect).

As this simple test case is not stationary, the correlation matrix is not Toeplitz circulant, as enforced in the SPOD proposed by Sieber *et al.* (2016). On the contrary, its localized patterns are the footprint of the different phenomena in the dataset and give to the POD the time localization capabilities that are not achievable by DMD/DFT.

The temporal correlation matrix of this synthetic test case, however, is almost defective, because it is close to having repeated eigenvalues. The three detected singular values are $\boldsymbol{\Sigma}_{\mathcal{P}} = \mathrm{diag}(297, 290, 284)$: although the energy contribution of the introduced mode is identical, the non-orthogonality of the temporal structures requires a slight energy re-distribution. Yet, the singular values are sufficiently close to the limit $\sigma_1 = \sigma_2 = \sigma_3$ to approach the condition of wholly undetermined POD.

The eigenvectors of $\boldsymbol{K}$, computed using the Matlab commands `eig` are shown in figure 9. These modes are a linear combination of the introduced ones, with a severe mixing of different phenomena over different modes. To further highlight the uniqueness problem, figure 10 shows that the temporal evolutions in figure 4 are almost eigenvectors of $\boldsymbol{K}$: the matrix multiplication $\boldsymbol{KT}_k$ leads to vectors that are close to simple multiples of the temporal evolution $\boldsymbol{T}_k$. The more $\boldsymbol{K}$ is close to being defective, the larger the number of vectors that are almost its eigenvectors and thus POD structures.

It is therefore clear that it is possible to obtain an orthonormal basis which is more representative of the coherent patterns in the data by constraining the POD within different frequency intervals (scales). In the proposed mPOD, this is achieved via the MRA of the correlation matrix, splitting the correlation matrix in $M = 4$ contributions with $\boldsymbol{F}_V = [0.4, 10, 20]$. The DFT of each scale is shown in figure 11(*a*); the dominant eigenvector for each scale is shown in figure 11(*b*). The spectral separation of the three contributions is achieved and the temporal structures recover the evolution of each mode. As a result, the spatial structures are identical to the introduced ones and are therefore not shown. At the cost of a minor loss in the decomposition

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

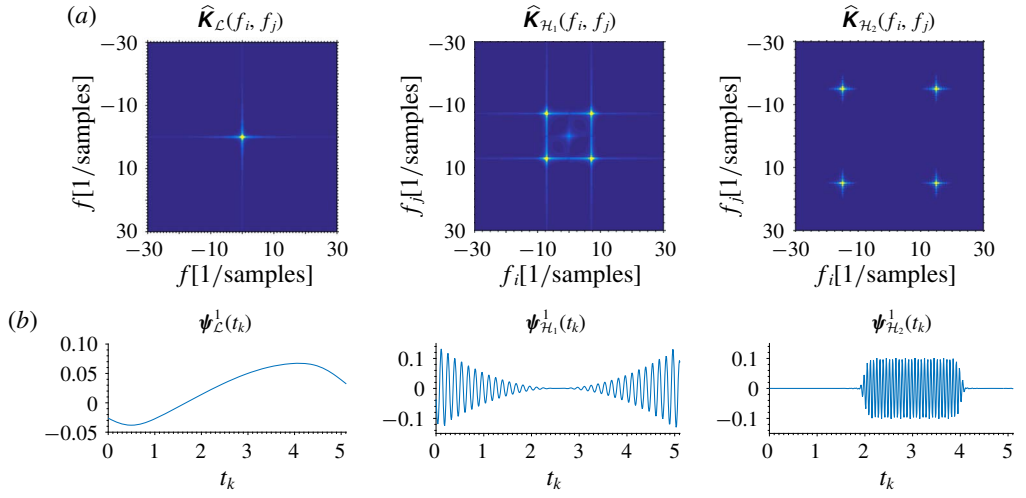


FIGURE 11. (Colour online) Fourier spectra of the correlation matrices $\boldsymbol{K}_{\mathcal{L}_1}$, $\boldsymbol{K}_{\mathcal{H}_1}$ and $\boldsymbol{K}_{\mathcal{H}_2}$ for the three scales identified via filter bank in the synthetic test case. (*b*) Shows the dominant eigenvector for each scale contribution. These are then assembled to construct the temporal structures $\boldsymbol{\psi}_{\mathcal{M}}$ as described in § 3.4.

convergence, if compared to the POD, the mPOD correctly identifies the coherent structures introduced and assigns a different mode to each of them.

5. Example II: source detection in nonlinear vortex dynamics

5.1. *Dataset description*

As a second test case, we propose a complex dataset which features the main physical mechanisms of complex fluid flows (i.e. nonlinear advection and diffusion), yet allowing for sufficient control over its dynamics so as to validate the decompositions. This test case is the numerical simulation of an advection–diffusion problem with random, pseudo-random and coherent sources. The problem is derived from the incompressible 2-D Navier–Stokes (NS) equations in the vorticity–streamfunction formulation. This formulation is common in the study of large-scale geophysical systems (Mariotti, Legras & Dritschel 1994; Flór, van Heijst & Delfos 1995; Trieling, van Wsenbeeck & van Heijst 1998), and consists of a set of two equations. The first is the transport equation of the vorticity $\omega = \boldsymbol{\nabla} \times \boldsymbol{u}$, where the velocity field $\boldsymbol{u} = (u, v)$ is described by a streamfunction ζ, defined so that $\partial_y \zeta = u$ and $\partial_x \zeta = -v$:

$$\partial_t \omega = -\mathcal{J}(\omega, \zeta) + \frac{1}{Re} \nabla^2 \omega + S_\omega, \tag{5.1}$$

where $\mathcal{J}(\omega, \zeta) = \partial_y \zeta \, \partial_x \omega - \partial_x \zeta \, \partial_y \omega$ is the Jacobian determinant of the solution vector (ω, ζ); Re is the Reynolds number weighting the importance of advection and diffusion; S_ω is the spatially and temporally varying source or sink of vorticity. The second equation is the Poisson problem linking vorticity and stream function:

$$\omega = \boldsymbol{\nabla} \times \boldsymbol{u} = \partial_x v - \partial_y u = -\nabla^2 \zeta. \tag{5.2}$$

Sources of vorticity S_ω in (5.1) are typically the vortex stretching/tilting due to the velocity gradient tensor, baroclinic effects due to the misalignment between pressure

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

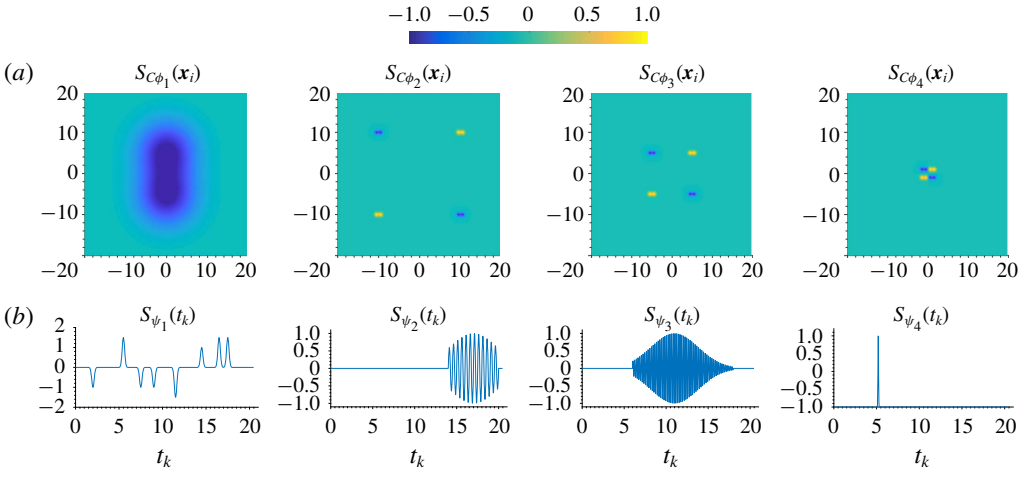


FIGURE 12. (Colour online) Spatial ($S_{C_{\phi_l}}$, *a*) and temporal ($S_{C_{\psi_l}}$, *b*) structures of the coherent source term $S_{C\omega}$ in (5.3).

and density fields and non-conservative body forces (e.g. the Coriolis effect). For the purpose of this work, we limit our interest to the nonlinear evolution of prescribed sources (regardless of their origin) and to the ability of data decomposition to correctly identify their spatial and temporal structures.

The solution of (5.1)–(5.2) is computed using a standard finite difference approach (Hoffmann & Chiang 1993; Kutz 2013), combining a fast Poisson solver for (5.2) and a Runge–Kutta time marching scheme for (5.1). The computational domain consists of a square Cartesian grid of $n_s = 128 \times 128$ points over a domain $\boldsymbol{x}_i \in [-20, 20] \times [-20, 20]$, while the solution vectors is exported in a time domain $t_k \in [0, 20.48]$ with a total of $n_t = 2048$ sample points. Periodic boundary conditions are set in all the boundaries.

The source term is treated explicitly, and consists of coherent ($S_{C\omega}$), pseudo-random ($S_{R\omega}$) and fully random ($S_{\mathcal{N}\omega}$) contributions:

$$S_\omega(\boldsymbol{x}_i, t_k) = \overbrace{\sum_{l=1}^{4} A_{Cl} S_{C_{\phi_l}}(\boldsymbol{x}_i) S_{C_{\psi_l}}(t_k)}^{S_{C\omega}} + \overbrace{A_R S_{R\phi}(\boldsymbol{x}_i) S_{R\psi}(t_k)}^{S_{R\omega}} + \overbrace{A_{\mathcal{N}} \mathcal{N}(\boldsymbol{x}_i, t_k)}^{S_{\mathcal{N}\omega}}. \tag{5.3}$$

The four spatial $S_{C_{\phi_l}}(\boldsymbol{x}_i)$ and temporal $S_{C_{\psi_l}}(t_k)$ structures of the coherent source $S_{C\omega}$ are shown in figure 12. The spatial structures are all constructed as sets of Gaussian vortex pairs of different size and position. The general form is

$$S_{C_{\phi_l}}(\boldsymbol{x}_i) = \pm \exp\left(\frac{(\boldsymbol{x}_i \pm \boldsymbol{x}_{Cl} \pm \boldsymbol{d}_l)^2}{\sigma_l^2} \right), \tag{5.4}$$

where the distance between the vortices in each pair are $\boldsymbol{d}_l = (d_{xl}, d_{yl}) = ([0, 0.5, 0.5, 0.5], [5.5, 0, 0, 0])$, the pair locations are $\boldsymbol{x}_{Cl} = (\boldsymbol{x}_{Cl}, \boldsymbol{x}_{Cl})$ with $\boldsymbol{x}_{Cl} = [0, 10, 5, 1]$ and the spreading factors are $\sigma_l = [49, 0.5, 0.5, 0.5]$.

The first source is composed of a single vortex pair of equal sign, which naturally merges producing a large recirculation, with angular velocity controlled by its inertia

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

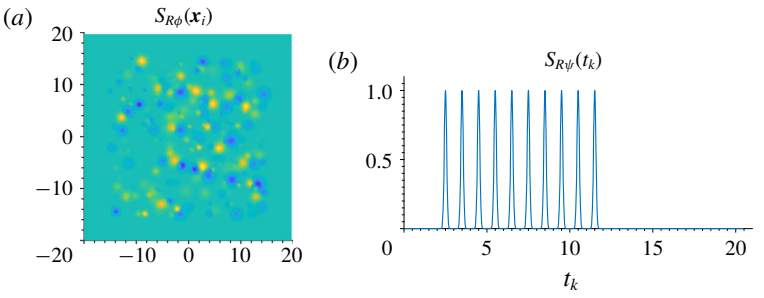


FIGURE 13. (Colour online) Exemplary spatial and temporal structures of the pseudo-random source term in (5.3). At each of the time pulse, a random distribution of 300 Gaussian vortices is introduced.

and the sign of the associated temporal structure $S_{C_{\psi_1}}$. This temporal structure is composed of a set of Gaussian pulses $A_i \exp((t_k - t_i)^2/\sigma_{t1}^2)$ of amplitude varying between $\pm[1, 1.5]$ and equal widths of $\sigma_{t1} = 0.2$. The large size of the vortex gives too much inertia to this structure to follow such sharp variations, but the change of sign in $S_{C_{\psi_1}}$ (for $t > 12$) changes the large-scale rotation from counter-clockwise to clockwise towards the end of the simulation. This process is much slower than all the other co-occurring.

The second and third coherent sources differ, in space, only in the distance between the vortex dipoles and have temporal structures consisting of Gaussian modulated harmonics:

$$S_{C_{\psi_{2,3}}} = \sin[2\pi f_{2,3}(t_k - \alpha_{2,3})] \exp\left[\frac{-(t_k - \beta_{2,3})^2}{\sigma_{t2,3}^2}\right], \tag{5.5}$$

with frequencies $f_{2,3} = [2, 5]$, equal width $\sigma_{t2,3} = 4$ but different phase delay for both the harmonic ($\alpha_{2,3} = [14, 6]$) and the modulation terms ($\beta_{2,3} = [17, 11]$). These differences in the phase delay create sharp variations that produce a regular vortex shedding. Finally, the fourth contribution is spatially similar to the second and third, with dipoles placed much closer to the centre of the domain, and a temporal structure consisting of a single Gaussian impulse located a $t_k = 5.2$ and width $\sigma_{t4} = 0.05$.

The pseudo-random source term $S_{R\omega}$ in (5.3) produces a periodic injection of random vortices. The spatial structure $S_{R\phi}$ consist of a set of 300 Gaussian vortices, randomly located with widths in the range [0, 1], and amplitudes in the range $[-1, 1]$. An exemplary set of vortices representing one of these structures is shown in figure 13(*a*). The temporal evolution $S_{R\psi}$, shown in figure 13(*b*), consists of 10 equal Gaussian impulses with width 0.1 and unitary amplitude and constant time shift. As the introduced spatial structures change randomly from one pulse to the other, this source contribution creates a chaotic background which has no spatial coherence but exhibits a temporal regularity. Finally, the white noise term $S_{\mathcal{N}_\omega}$ in (5.3) is introduced in the post-processing step, using a random number generator varying between 0 and 1 to generate the spatio-temporal noise $\mathcal{N}(\boldsymbol{x}_i, t_k)$.

The coefficients A_{C_l} ($l \in [1, 4]$), A_R and $A_\mathcal{N}$ in (5.3) set the relative importance of these three contributions and are taken as $A_l = [9, 610, 1100, 2500]$, $A_R = 300$ and $A_\mathcal{N} = 7$. These coefficients were estimated, by trial and error, to give approximately the same importance to the four coherent sources, taking into account the strong damping effect produced by diffusion (especially on the smaller and faster scales), obtained by

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

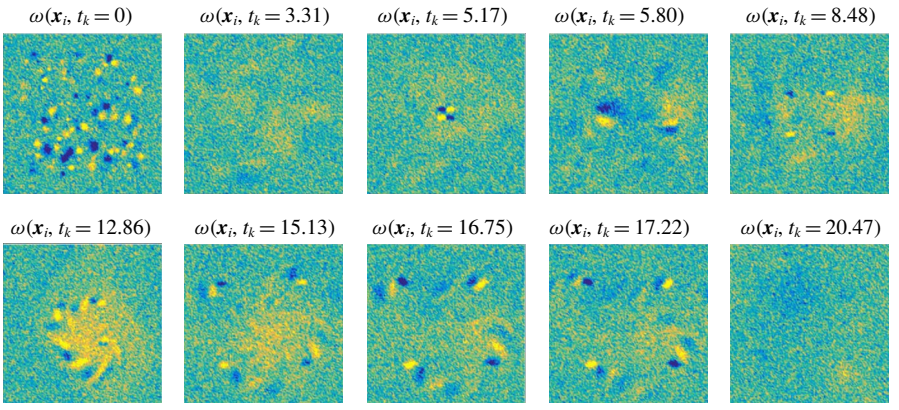


FIGURE 14. (Colour online) Exemplary snapshots from the advection–diffusion problem in (5.1). These selected snapshots describe some important instants in the evolution of the vorticity field considered in this test case. An animation of this test case is available as supplemental Movie 2.

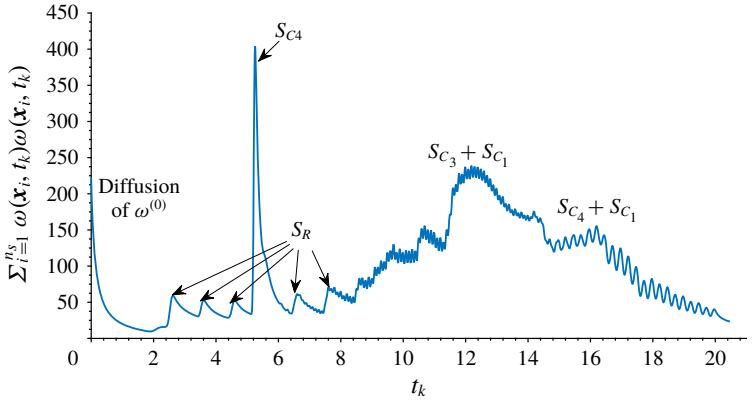


FIGURE 15. (Colour online) Time evolution of the global energy content (norm of each vorticity realization), extracted from the diagonal of the temporal correlation matrix $\boldsymbol{K}$.

setting the Reynolds number in (5.1) to $Re = 1$. The initial vorticity field is a random set of vortices similar to those introduced by the pseudo-random source term.

To visualize the complexity of the dataset obtained, figure 14 collects 10 selected snapshots of the vorticity field, each corresponding to the time indicated in the title. The description of the vorticity evolution in these snapshots is well complemented by the evolution of the energy signal in time, that is the norm of each vorticity snapshot in the diagonal of the temporal correlation matrix K. This signal is shown in figure 15, and the occurrence of some key events are indicated with an annotation.

For $t_k < 2$, the dataset describes the chaotic interaction of the initial field, which strongly diffuses to almost vanish (cf. figure 15). Before this occurs, the first coherent source gives its first pulse, producing a large-scale counter-clockwise rotation in the background. Until $t_k = 5$, the three injections of random vortices produces three noticeable increase in the energy content against diffusion, but no coherent patterns can be identified at this stage.

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

At $t_k \approx 5.15$, the injection of the source term $S_{C\psi_4}$ stands out from the chaotic motion, and the energy in the vorticity field reaches its highest peak. The four vortex pairs immediately merge into two identical dipoles of equal vorticity, which repel from the centre of the domain in opposite direction and speed. However, due to the background rotation produced by the first source and the chaotic interaction with the other random injections from the pseudo-random term $S_{R\omega}$, these two dipoles do not travel the same distance and are broken and wholly dissipated by $t \approx 6.5$.

At $t_k > 6$, the third coherent source begins its pulsation. At $t_k \approx 7$ this generates four wakes which shed at a constant frequency. The overall vorticity in the domain increases and the large-scale clockwise rotation involves the four wakes in a closed circular pattern. Before this large-scale rotation vanishes, at $t_k \approx 14$, the second coherent source starts to pulse at a lower frequency, producing other four trains of vortices, initially also rotating counter-clockwise. Until approximately $t_k = 17$ the second and third source coexist and produces patterns of similar energy content but different frequencies.

From $t_k > 17$, the second source vanishes and the vortex pair introduced by the last three pulses of the first coherent source slowly changes the direction of the rotation. This breaks the circular shedding into two vertical ones before a new circular shedding rotating in the opposite direction is produced. The change of direction and the strong diffusion reduce the energy content of the flow until only a large-scale vortex, rotating clockwise, slowly dissipates.

5.2. *Modal analysis*

The relative L^2 error convergence (4.2) for the DFT, cDMD, sDMD, POD and mPOD are shown in figure 16. The cDMD and the sDMD formulations yield very similar (yet not identical) results if the full POD basis is used in the construction of the reduced propagator in (2.19).

This test case strongly highlights the convergence limits of the DMD for noisy data with a nonlinear evolution. As shown in § 2.3, the amplitude coefficients $\sigma_{\mathcal{D}r}$ in the DMD are projections of the initial conditions onto the DMD basis: large amplitudes are assigned to modes which are more strongly linked to the initial conditions, even if these vanish immediately afterwards – in this case due to the strong diffusion. As a result, both algorithms produce DMD eigenvalues lying well inside the unit circle and exponentially decaying. Since decaying modes with real frequencies are not orthogonal, the DMD convergence is drastically compromised and yields, for $r_c < 200$, an L^2 approximation error which is larger than the norm of the dataset itself.

The problem is particularly evident for the sDMD with the reduced propagator $\boldsymbol{\mathcal{S}}$ expressed in (2.19). Figure 17 shows the eigenvalues distribution in the complex plane for reduced propagators $\boldsymbol{\mathcal{S}}$ of different dimensions $r_c \times r_c$, obtained by projecting the full propagator onto the first r_c POD modes. For $r_c \ll n_t$, most of the eigenvalues are well inside the unit circle, and their contribution to the decomposition is limited to the first few time steps. Since each of these modes is normalized, the shortest is their duration, the largest the associated amplitude. As r_c approaches n_t and the DMD basis approaches completeness, the number of modes with exponential decays is reduced, and the sDMD approaches the cDMD.

The amplitudes of the modes for the DFT and the cDMD are mapped in the complex plane in figure 18. The DFT spectrum shows that the energy contribution of the mean flow is stronger than in the previous test case, and highlights three dominant

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

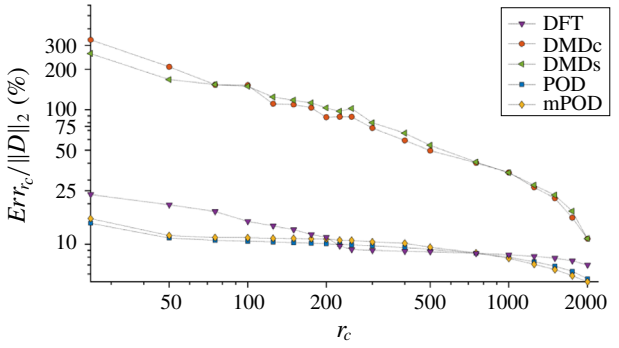


FIGURE 16. (Colour online) Decomposition convergence for DFT, cDMD, sDMD, POD and mPOD for the vorticity field in the nonlinear advection–diffusion problem considered in this section.

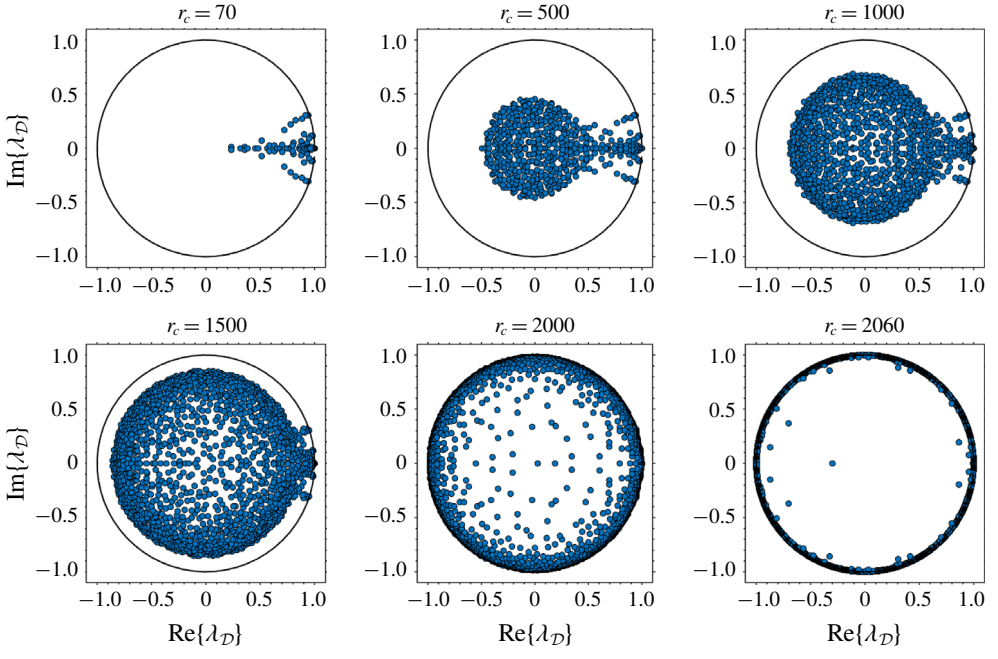


FIGURE 17. (Colour online) Eigenvalue distribution for reduced propagators $\boldsymbol{\mathcal{S}} \in \mathbb{C}^{r_c \times r_c}$ in (2.19) obtained by projecting the full propagator $\boldsymbol{P}$ onto the first r_c POD modes. The smaller the propagator, the more the decomposition convergence is limited within a smaller number of time steps from the initial data. At $r_c = n_t - 1$, the sDMD approaches the cDMD.

frequencies together with a wide range of modes with nearly identical energy content. These are produced by the high random noise level in the dataset and the presence of impulsive events. For the cDMD, the random noise makes the calculation of the companion matrix more difficult, producing an error of 92 % of the L^2 norm of the last realization in the minimization $\|\boldsymbol{d}_{n_t} - \boldsymbol{D}_1\boldsymbol{c}\|_2$.

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

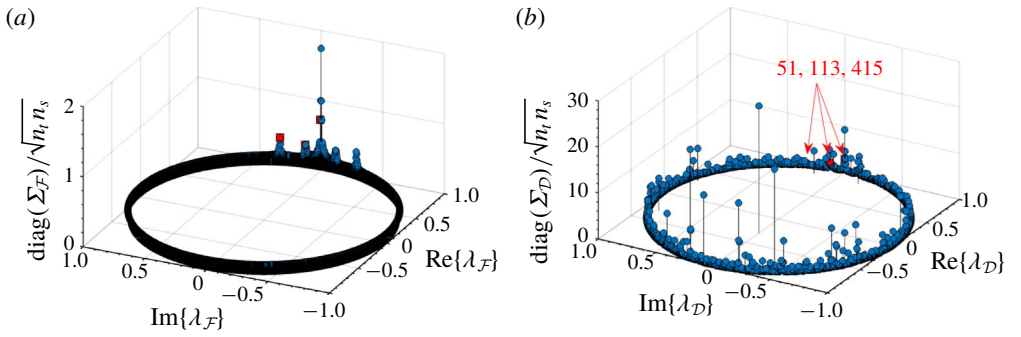


FIGURE 18. (Colour online) Amplitude spectra of the DFT (*a*), and cDMD (*b*) mapped in the complex plane. While both frequency based decompositions are symmetric along the imaginary axis, the DMD is sensitive to the high noise level and the strong diffusion of the initial dataset, which results in all the leading modes lying inside the unit circle.

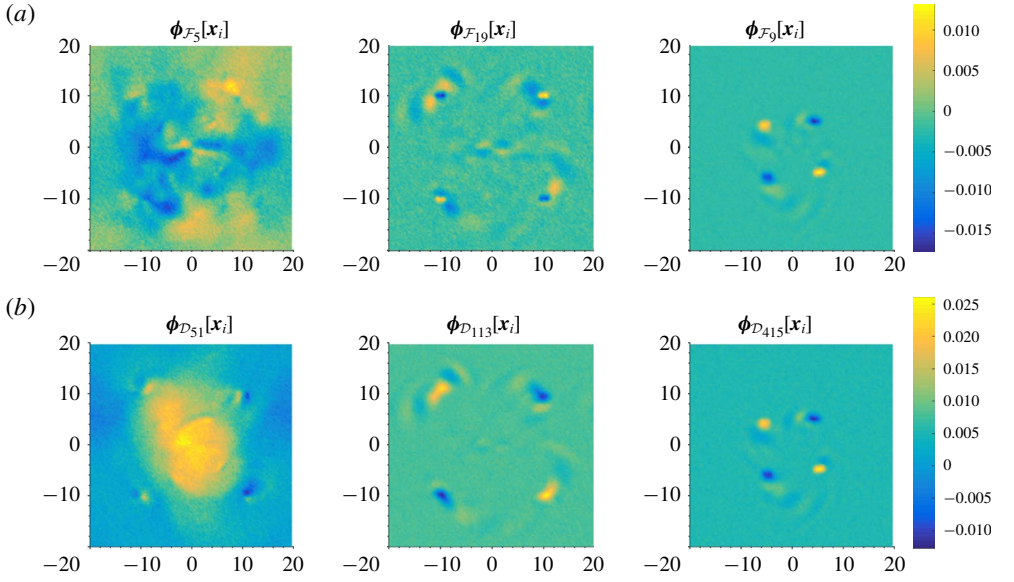


FIGURE 19. (Colour online) (*a*) Spatial structures $\boldsymbol{\phi}_{\mathcal{F}}[\boldsymbol{x}_i]$ of three dominant DFT modes, pulsing, from left to right, at frequencies of $f_n = [0.12, 2.01, 5.02]$. These corresponds to leading frequencies of the coherent source term. (*b*) Spatial structures $\boldsymbol{\phi}_{\mathcal{D}}[\boldsymbol{x}_i]$ of the dominant sDMD modes with frequencies [0.08, 1.97, 4.99], that is the ones closer to those dominating the DFT spectra.

The three dominant frequencies in the DFT are $f_n = [0.12, 2.01, 5.02]$; the first is linked to the large-scale motion produced by the first coherent source, the second and third are related to the frequencies of the second and third sources pulsing at $f_{2,3} = [2, 5]$. Figure 19 (*a*) shows the corresponding spatial structures, which are well related to the introduced ones. The three cDMD modes with unitary modulus having frequencies closer to these DFT modes – also captured by all the sDMD in figure 17 – are shown in figure 19(*b*). The amplitude of these modes is labelled with a red

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

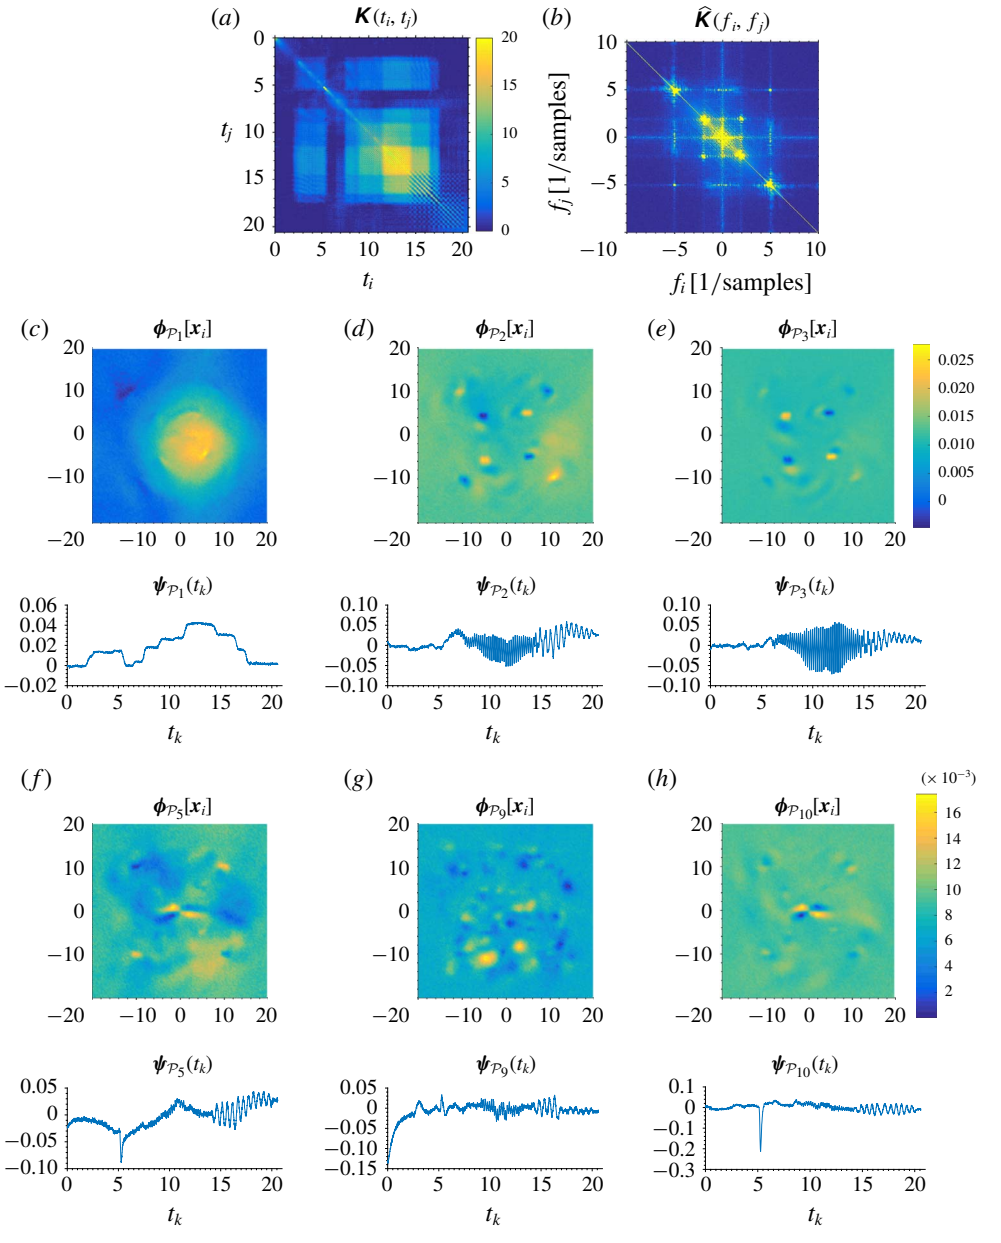


FIGURE 20. (Colour online) Same as figure 9 but for the second test case discussed in this section.

marker in the complex plane of figure 18, and is much lower than the dominant ones, associated with the vanishing initial conditions.

As for the synthetic test case in § 4, both the DFT and the DMD have sufficient frequency resolution to correctly localize periodic phenomena in the frequency domain and – thanks to the projection in (2.6) – in the space domain. However, as no temporal localization is possible, the convergence of these decompositions is remarkably poorer

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

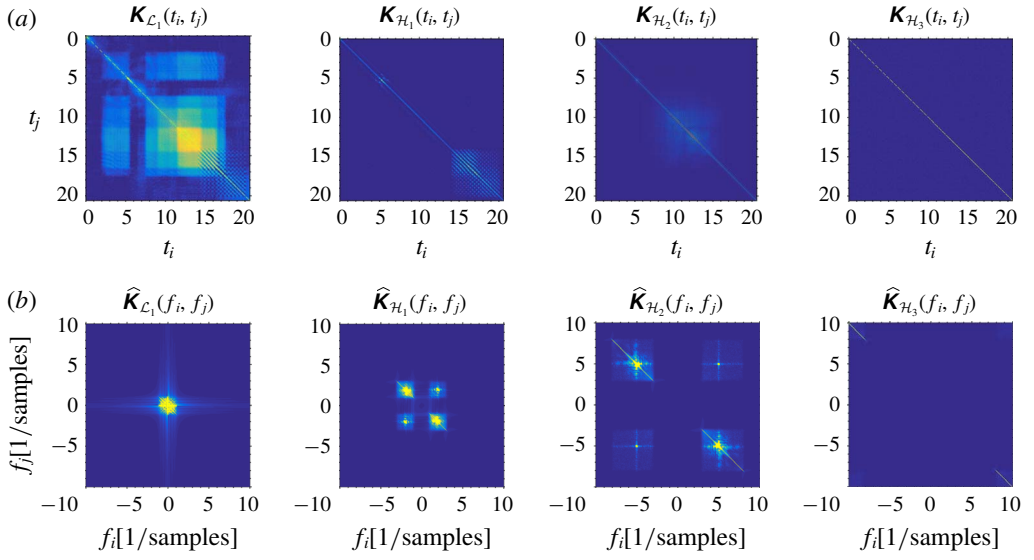


FIGURE 21. (Colour online) Results from the MRA of the correlation matrix $\boldsymbol{K}$, split into four contributions by a filter bank with frequency splitting vector $\boldsymbol{F}_V = [1, 3, 8]$. (*a*) Shows the temporal correlations; (*b*) shows their frequency content.

than the POD and results in a high redundancy of the spatial basis. This is particularly evident in the description of the impulsive event produced by the fourth coherent source, the footprint of which appears in a large number of DFT and DMD modes.

The POD has complementary advantages and limitations. Figure 20 shows the results of the POD analysis, including the temporal correlation matrix $\boldsymbol{K}$ and its DFT spectra $\widehat{\boldsymbol{K}}$, together with the spatial and the temporal structures of six representative POD modes.

The temporal correlation matrix $\boldsymbol{K}$ gives information on the time localization of different events, featuring two significant chessboard patterns corresponding to the harmonic contributions of the second and third sources shown in figure 12, and rectangular areas corresponding to the time interval of major activity of the first large-scale coherent source. In the diagonal, a clear peak is located at $t_k = 5.2$, corresponding to the introduction of the fourth impulsive source, while the random noise gives a pedestal correlation level along the whole diagonal.

The first mode captures the large-scale motion produced by the first coherent source, with a temporal structure slightly polluted by white noise. Uncorrelated noise is uniformly distributed over the entire POD spectra (Mendez *et al.* 2017) and the low noise level present in this mode is due to the large dataset available. The second and third modes show a severe spectral mixing, as their energy contribution is similar ($\sigma_{\mathcal{P}_2} = 4272$ and $\sigma_{\mathcal{P}_2} = 4433$). This problem, also illustrated on the simple synthetic test case in § 4, results in a combination of the features from the second and third coherent sources. Such mixing extends to other POD modes with different phase delays, to cancel with the summation of multiple modes. The fifth mode mixes three phenomena: a large-scale contribution to the rotation produced by the first source, a harmonic contribution due to the second source and a sharp peak due to the impulsive fourth source. The same occurs with the ninth mode, which mostly describes the diffusion of the initial data. Interestingly, the tenth mode well captures

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

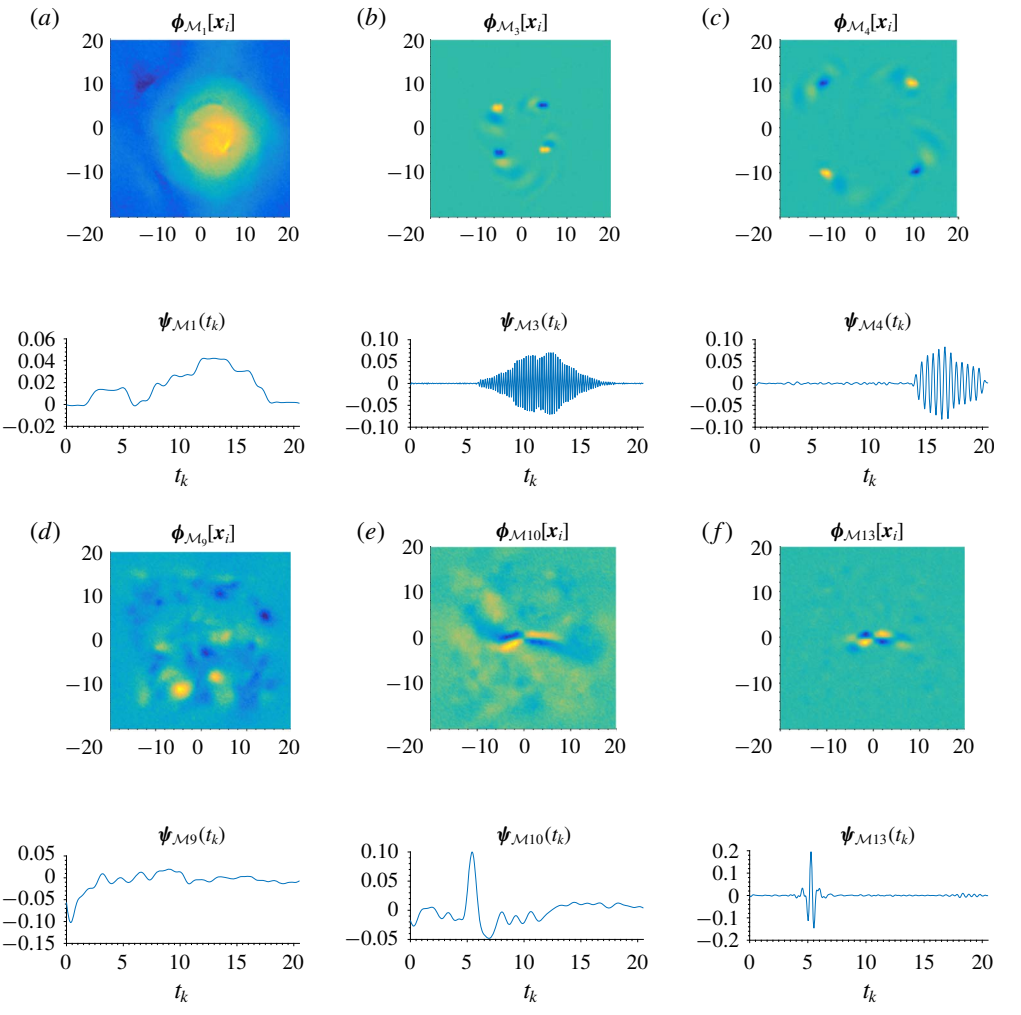


FIGURE 22. (Colour online) Spatial and temporal structures for six representative mPOD modes for the nonlinear vortex test case presented in this section.

the impulsive event, although both the harmonic contribution of the second and third sources are present to a minor extent.

Figure 21 shows the results of the MRA analysis of the temporal correlation matrix $\boldsymbol{K}$, using $\boldsymbol{F}_V = [1, 3, 8]$ and computing the filter orders from (A 5). (*a*) Displays the four contributions of the identified scales and (*b*) their DFT. No significant phenomenon appears in the last contributions, which mostly contains a substantial portion of the noise energy, whereas the other three scales isolate the contributions of different sources. Being the overlapping of frequencies negligible, the eigenvectors of these matrices are almost orthogonal complements for $\mathbb{R}^{n_t}$, and the resulting full basis $\boldsymbol{\Psi}^0_{\mathcal{M}}$ (3.15) requires a minor re-orthogonalization via $\boldsymbol{QR}$ factorization.

The spatial and temporal structures of six representative mPOD modes are shown in figure 22. These modes are similar to the POD modes in figure 20, but with no frequency overlapping. The first mode is a cleaner version of the first POD mode, with negligible random noise; the third and fourth capture the contributions of the second

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

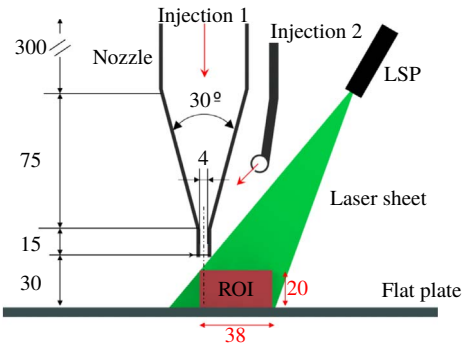


FIGURE 23. (Colour online) Schematic of the experimental set-up for TR-PIV analysis of an impinging plane jet, showing the region of interest (ROI). The dimensions are in mm.

and third coherent sources with no mixing. The sixth mode describes the diffusion of the initial data while the tenth and the thirteenth are related to the impulsive event. The frequency bounding makes the decomposition stronger than the standard POD versus the random noise since only a minor portion of its contribution remains in each scale. Yet, the capabilities of capturing impulsive phenomena (also distributed over the entire frequency spectra) are reasonably preserved: traces of this impulse event are present in all the scales (the tenth mode is the eigenvector of $\boldsymbol{K}_{\mathcal{L}}$; the thirteenth is eigenvector of $\boldsymbol{K}_{\mathcal{H}_3}$), but no significant overlapping between this event and the other harmonics is produced.

6. Example III: particle image velocimetry of an impinging jet flow

6.1. *Dataset description*

As a third test case, we propose an experimental dataset obtained via time-resolved particle image velocimetry (TR-PIV) on a classical, statistically stationary flow configuration: a plane jet impinging normally onto a flat plate. Impinging jet flows have received considerable attention for their application in drying, cooling or heating processes (Zuckerman & Lior 2006; Buchlin 2011; Specht 2014) and the POD of PIV data has been successfully used by many authors to identify the resulting coherent structures (Kim *et al.* 2007; Hammad & Milanovic 2009; Charmiyan *et al.* 2017; Pieris *et al.* 2017).

These flows are characterized by largely different time scales in three main regions: the free jet, close to the nozzle exit; the stagnation region, close to the flow impact; the wall jets, released parallel to the wall from impact (Gutmark, Wolfshtein & Wygnanski 1978).

The experimental set-up is shown in figure 23. The rectangular outlet section of the nozzle has an opening of $H = 4$ mm and an aspect ratio of $H/W = 62.5$. The chamber gauge pressure P_N is monitored by a piezoresistive pressure transducer AMS5812 connected to a pressure tap and the flat surface consist of a polished aluminium plate placed at a distance of $Z = 30 \pm 1\,\%$ mm. The experiment is carried out with $P_N = 30 \pm 1\,\%$ Pa leading to an average jet flow velocity of $U_0 = 6.5 \pm 2\,\%$ m s^{-1}, that is a Reynolds number of $Re \approx U_0 H/\nu = 1733$ considering a kinematic viscosity of $\nu = 15$ cSt.

The seeding tracers for the PIV measurements are microdroplets of mineral oil (Ondina Shell 917), produced by a Laskin nozzle (PIVTEC45-M). These are

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

introduced both in the stagnation chamber and in a second injection on one side of the nozzle to seed the entrainment flow. The laser source is a dual diode-pumped ND:YLF laser (Quantronix Darwin Duo 527), reaching the test section via a light sheet probe (LSP, from ILA). The videos are acquired by a CMOS Photron FASTCAM SA1, positioned at about 30 cm and equipped with a zoom objective Tamron CZ-735 70–300 mm to provide an image scaling factor of 24.1 ± 0.1 pixel mm^{-1}.

The investigated region of interest (ROI) is a rectangular area of 20 mm $\times$ 38 mm at a distance of 10 mm from the nozzle outlet (cf. figure 23). The video sequence is made of $n_t = 2728$ images pairs of 1000×740 pixels, acquired in frame-straddling mode at $f_s = 2$ kHz. The separation time between the frames is set to $\Delta t_p = 70$ µs to produce a maximum displacement, in the jet flow, of approximately 8 pixels. The PIV evaluation is carried out with the Matlab package PIVlab (Thielicke & Stamhuis 2014), using standard iterative multi-step interrogation (four passes from 128×128 to 8×8 windows) with spline window deformation, 2-D Gaussian sub-pixel interpolation and vector validation via a median test. The resulting vector field consists of $n_s = 60 \times 114$ spatial realizations sampled on a uniform Cartesian grid.

An exemplary instantaneous flow field is shown in figure 24, together with the location of six probes. The normalized power spectral density computed from the velocity magnitude in these locations is shown in figure 24(*b*) as a function of the dimensionless frequency $St = f_n H/U_0$ and normalized by the corresponding maximum value. These spectra are computed using the Welch's method (Welch 1967) with Hanning windows of $n_t/4$ width and 25 % overlapping. The first probe P_1 is located in the free-jet flow, where the shear layer instabilities develop unaffected by the presence of the impinged wall. At this location, most of the energy is centred around a peak at $St \approx 0.3$, in agreement with other studies on the free plane jet in the transitional regime (Thomas & Goldschmidt 1986; Suresh *et al.* 2008). This corresponds to the formation of large vortices (primary vortices) which originates from a Kelvin–Helmholtz (K–H) instability and evolve into roller-like structures (Mumford 1982).

As the flow approaches the stagnation region (probes P_2, P_3), the frequency spectra broaden and moves towards lower frequencies due to the vortex merging and the velocity decay, as also reported by Pieris *et al.* (2017) for impinging jet flows at higher Reynolds numbers. Far from the stagnation region, the wall jet flow (P_4, P_5, P_6) is governed by a much lower frequency and the unsteady separation due to the intermittent passage of the primary vortices (Didden & Ho 1985; Hadžiabdič & Hanjalič 2008).

6.2. *Modal analysis*

Figure 25 shows the L^2 convergence in (4.2) for POD, cDMD, DFT and two mPOD with frequency splitting vectors. The first, in terms of Strouhal numbers, consists of $M = 4$ with $\boldsymbol{F}_{V1}(m) = [0.1, 0.2, 0.5]$; the second consists of $M = 50$ with $\boldsymbol{F}_{V2}(m) = 0.01 + m\Delta f$ with equal bandwidth $\Delta f = 0.02$. As discussed in § 3.2, the finer the frequency splitting, the more the mPOD approaches a DFT.

The frequency spectra for DFT and cDMD are mapped in figure 26. The investigated jet flow is in statistically stationary conditions and the mean flow accounts for a significant portion of the dataset. For plotting purposes, therefore, the entry corresponding to the zero frequency is removed. The spectra of the DFT shows that the investigated test case is rather broadband, with few frequencies reasonably standing over the others. Despite the significant random noise and turbulence contribution,

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

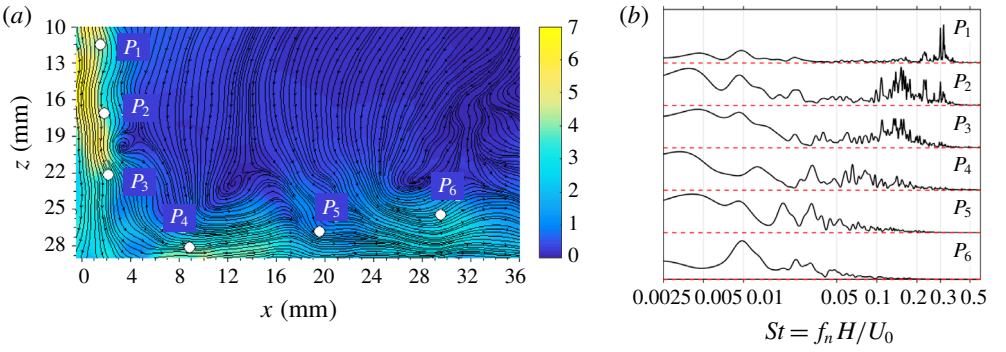


FIGURE 24. (Colour online) (*a*) Exemplary snapshot of velocity field from TR-PIV of an impinging jet flow, with contour plot displaying the magnitude of the velocity in m s^{-1}, and location of six representative probes. An animation of this test case is available as supplemental Movie 3. (*b*) Normalized power spectral densities in the probes in (*a*), versus Strouhal number.

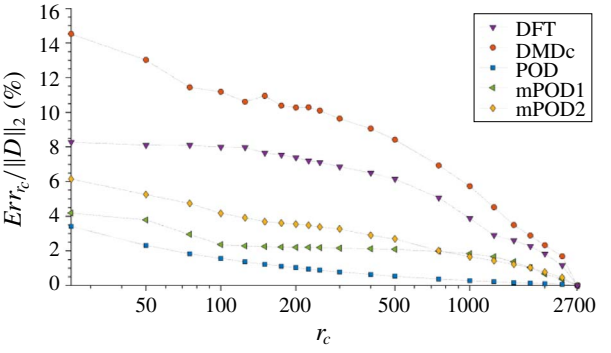


FIGURE 25. (Colour online) L^2 convergence (4.2) of the POD, cDMD, DFT and two mPOD with different frequency splitting vectors for the TR-PIV field of the impinging jet flow.

the presence of a statistically converged time average results in a relative error $\|\boldsymbol{d}_{n_t} - \boldsymbol{D}_1\boldsymbol{c}\|_2/\|\boldsymbol{d}_{n_t}\|_2 \approx 6\,\%$ in the minimization defining the companion matrix for the cDMD. The resulting DMD dominant modes identify the relevant frequencies more sharply than the DFT. Nevertheless, several eigenvalues lie inside the unit circle and the convergence of the cDMD, although not dramatically poor as in the previous test case, is below the DFT.

The spatial structures of two representative DFT and DMD modes are shown in figure 27. For the DFT, these structures are associated with dimensionless frequencies of $St = [0.033, 0.299]$ and are thus expected to be linked to the large-scale oscillations of the wall jet and K–H structures produced in the free-jet region respectively (see figure 24). The constraint of a single harmonic is nevertheless too restrictive to isolate these mechanisms, and the corresponding spatial modes are not capable of fully identifying the related coherent patterns. The DMD spatial structures, pulsing at dimensionless frequencies of $St = [0.009, 0.293]$, appear less noisy than the DFT ones, thanks to the better frequency selection. The mode $\boldsymbol{\Phi}_{\mathcal{D}_{27}}$ describes the large-scale

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

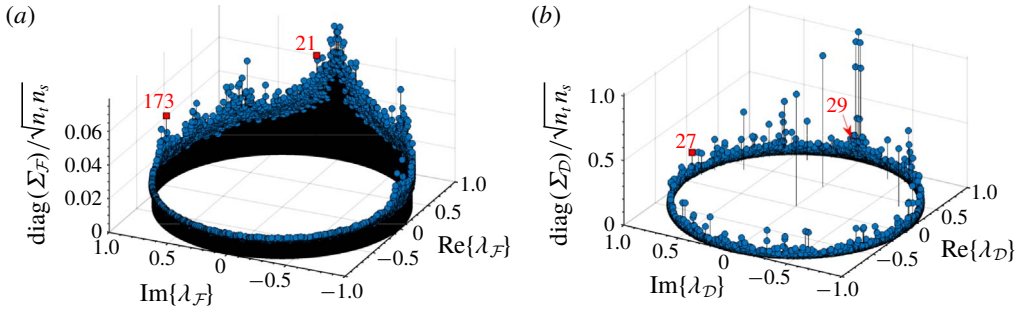


FIGURE 26. (Colour online) Same as figures 6 and 18, but for the third test case presented in this section.

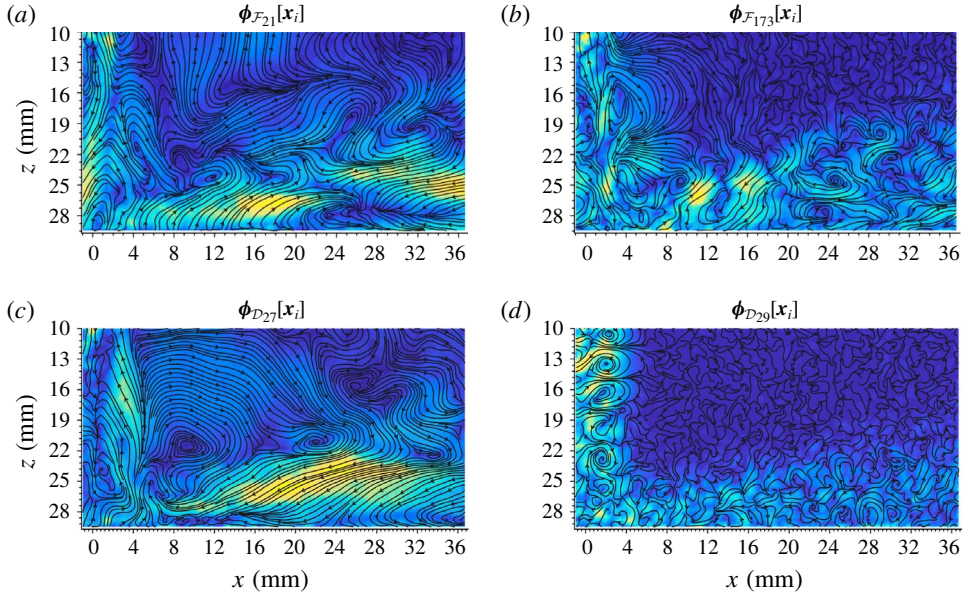


FIGURE 27. (Colour online) Spatial structures of two representative DFT (*a*,*b*) and DMD (*c*,*d*) modes, with amplitudes shown in figure 26 with a red marker. The DFT modes are associated with dimensionless frequencies of $St = [0.033, 0.299]$; for the DMD these are $St = [0.009, 0.293]$.

lifting of the wall jet, due to the interaction with the roller structures that in this region lose most of their momentum and mix with the entrainment flow. As expected, the mode $\boldsymbol{\Phi}_{\mathcal{D}_{29}}$ captures the K–H structures, vortical structures approximately 2 mm in diameter that arise from the primary shear layer instability of the free-jet flow.

The results of a standard POD analysis are collected in figure 28. Since the POD modes are stationary, it is interesting to show the frequency content $\widehat{\boldsymbol{\psi}}_{\mathcal{P}_r}$ associated with each spatial structure, rather than their temporal evolution $\boldsymbol{\psi}_{\mathcal{P}_r}$. All the modes in the range $r = 4$ to $r = 11$ are related to travelling structures and arise therefore in pairs that have $\pi/2$ phase delay in space and time. Only the four representative modes are shown.

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

The temporal correlation matrix $\boldsymbol{K}$ displays a regular pattern composed of a broad range of frequencies and $\widehat{\boldsymbol{K}}$ displays a band-like structure. Yet, a dominant frequency around 12 Hz ($St = 0.0074$) is visible. This corresponds to the large-scale flow motion in the wall jet flow, which covers a larger portion of the domain and, therefore, gives a more significant contribution to the overall correlation level. As the data are statistically stationary, the constant vector $\boldsymbol{\psi}_\mu = \underline{1}/\sqrt{n_t}$ is almost an eigenvector of the correlation matrix, and thus the first POD mode reproduces closely the mean flow (not shown). For plotting purposes, the mean correlation is removed from $\widehat{\boldsymbol{K}}$.

The spatial structures identified by the POD are considerably less influenced by the uncorrelated portion of the dataset and highlight coherent features in the flow. Besides the mean flow, the dominant contribution $\boldsymbol{\phi}_{\mathcal{P}_2}$ in figure 28 is mostly related to the large-scale pulsation in the wall jet (labelled with *I*) and its periodic separation, as expected from the spectra of the correlation matrix $\widehat{\boldsymbol{K}}$. However, this spatial structure also captures, to a minor extent, coherent structures in the free-jet region (labelled with *II*). These are expected to occur at much higher frequency (in the range $St = 0.1$–0.2) as shown by the probes P_2–P_3 in figure 24.

The frequency spectrum $\widehat{\boldsymbol{\psi}}_{\mathcal{P}_2}$ gives the final proof of spectral mixing for this POD mode: a minor energetic content is also present in the frequency range belonging to the flow structures of the free-jet flow, and the projection of the data along this temporal evolution propagates this mixing to the spatial domain. The same occurs in the mode $\boldsymbol{\phi}_{\mathcal{P}_3}$, which is mostly linked to the travelling vortices in the jet, but – to a minor extent – also to the wall jet separation. All the other modes suffer from the same problem, which reaches its worst case in the mode $\boldsymbol{\phi}_{\mathcal{P}_9}$. In addition to the two mechanisms previously described, this mode also captures the K–H structures produced further upstream, where they originate at $St \approx 0.3$ as clearly shown by the probe P_1 in figure 24. This test case shows that the spectral mixing problems of the POD, highlighted by the two test cases of §§ 4 and 5, can easily occur also in statistically stationary flows if these are governed by largely different scales with similar energy content.

As for the previous two test cases, the spectral separation produced by the MRA is sufficient to isolate the contribution of different scales at the cost of a minor reduction in the decomposition convergence (cf. figure 25). Four representative mPOD modes, including spatial structures and related frequency content, are shown in figure 29. These are constructed by setting $\boldsymbol{F}_{V1}(m) = [0.1, 0.2, 0.5]$, producing the L^2 convergence labelled as mPOD1 in figure 25. The resulting mPOD modes are very similar to the POD counterparts, showing that these spectral constraints have a minor impact on the decomposition. Yet, no spectral mixing is produced. The second mode captures only the dynamics of the separating wall jet, the fourth is related only to the evolution of the primary vortex advected towards the wall, while the ninth captures only the formation of the primary vortices in the free-jet region. The seventh mode describes the large-scale oscillation of the jet which is associated with the stagnation region and the initial stage of the wall jet development.

By increasing the number of frequency bandwidths, one could further narrow the frequency content of each mode, improving the frequency resolution but decreasing the L^2 convergence of the method as shown in figure 25 by the second mPOD with finer frequency splitting vector. This can alter the ranking of different modes, giving priority to mechanisms that have narrower frequency content, and eventually move the mPOD towards the DFT for unreasonably fine frequency splitting vectors.

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

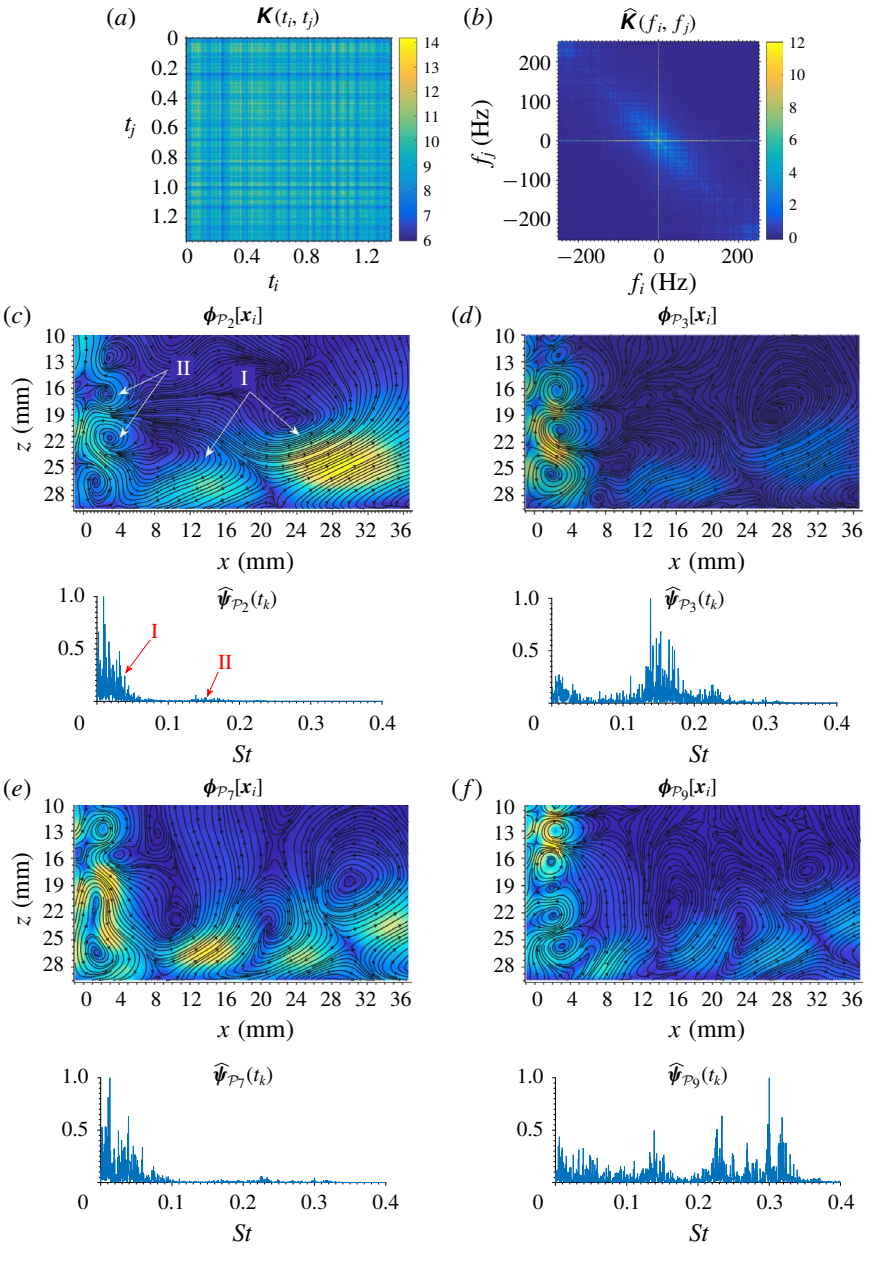


FIGURE 28. (Colour online) Same as figures 9 and 20, but for the third test case discussed in this section and replacing the temporal structures $\boldsymbol{\psi}_{\mathcal{P}}$ with their normalized frequency content $|\widehat{\boldsymbol{\psi}}_{\mathcal{P}}|/\max\{\widehat{\boldsymbol{\psi}}_{\mathcal{P}}\}$.

7. Conclusions

The state of the art of data-driven decomposition has been extensively discussed, with particular emphasis on the need to bridge the gap between the energy optimality of the proper orthogonal decomposition (POD) and the spectral purity of the dynamic

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

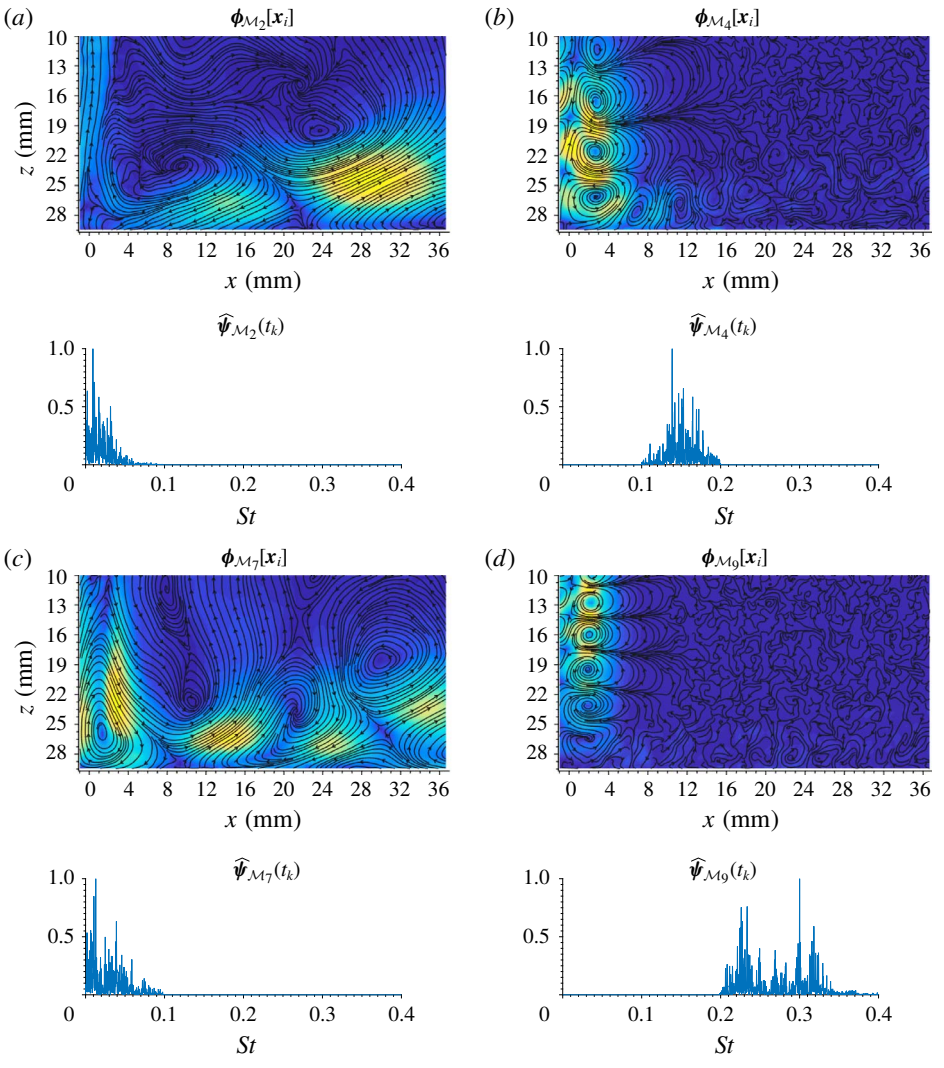


FIGURE 29. (Colour online) Selected mPOD modes corresponding to the POD modes in figure 28.

mode decomposition (DMD). A generalized algebraic framework is introduced, and all the investigated decompositions are presented as matrix factorizations. This framework is then used to derive a novel decomposition, referred to as multi-scale proper orthogonal decomposition (mPOD).

The derivation of the mPOD starts with an investigation on the link between the frequency content (the DFT) of the correlation matrix and that of its eigenvectors (the POD modes). This link is used to study the impact of a time filter in the dataset on the frequency of the correlation matrix and thus the related POD modes. Specifically, it is shown how a filter bank can be used to perform multi-resolution analysis (MRA) of the correlation matrix, splitting it into the contribution of different scales retaining non-overlapping portions of the frequency spectra of the data. These scales are

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

equipped with their optimal eigenbasis which can be kept mutually orthogonal and finally assembled into a single mPOD basis.

The novel decomposition is tested on three selected test cases and its performances compared to those of a discrete Fourier transform (DFT), a standard POD and the DMD in its two standard formulations: the Arnoldi-based approach (cDMD) and the SVD-based approach (sDMD).

The first of the proposed test case is a synthetic dataset with three predefined modes. These have equal energy, identical but shifted spatial structure and temporal evolution with different frequency and time localization. None of the popular decompositions presented proved capable of distinguishing the introduced modes. Harmonic decompositions such as DFT and DMD allow for proper frequency identification, but yield poor convergence, high redundancy in the spatial structures and a complete lack of temporal localization due to the infinite duration of the time basis. On the contrary, the POD allows for the best convergence, with only three modes required to describe correctly (i.e. with zero L^2 errors) the dataset. However, it produces a 'spectral mixing' between different modes. This mixing results in modes that are linear combinations of the expected ones since their similar energy content results in almost repeated eigenvalues of the correlation matrix and thus almost undefined (non-unique) POD.

The proposed mPOD proved capable of distinguishing the three introduced modes, with no spectral mixing, excellent time localization and almost optimal convergence.

The second investigated test case is a 2-D numerical simulation of the Navier–Stokes equation, with prescribed coherent, random and pseudo-random sources. The complexity of the data, further polluted with white noise, strongly accentuates all the limits highlighted by the previous test case, with the DMD L^2 relative error of the order of 200 % within the first modes and up to three different mechanisms sharing the same POD mode. Neither the DFT/DMD nor the POD are capable of identifying and localizing in space and time the different coherent source terms, while the mPOD succeeds at the scope with a negligible loss of convergence over the POD.

Finally, the third case is the time-resolved flow field of an impinging jet experimentally investigated via TR-PIV. Although the data are stationary and statistically well converged, the spectral mixing of the POD is observed while the harmonic constraints of the DFT and DMD yield much poorer convergence and noisy spatial structures. The mPOD naturally handles both problems.

To conclude, the discussed test cases prove that a minor spectral constraint in the diagonalization of the correlation matrix allows overcoming of the frequency limits of the POD, and show that the proposed mPOD provides an excellent compromise between energy optimality and spectral purity. While this work has focused on data-driven decompositions for fluid flows, the possibility of constraining the POD optimality within specific scales of interest gives to the mPOD more flexibility than POD or DMD/DFT for a large variety of applications. In data filtering/compression, for example, the mPOD is more robust against random noise (equally spread over all the POD/DFT spectra) and allows to better focus on particular features (e.g. allowing for different levels of data compression/filtering for different scales). This could find useful applications, for example, in the background/foreground separation problem in video analysis, where the slow motion of large objects is modelled by few dominant POD modes (or slow DMD ones) while the dynamics of smaller objects is captured by the remaining modes (see Oliver, Rosario & Pentland 2000; Grosek & Kutz 2014; Sobral *et al.* 2015; Mendez *et al.* 2017). The mPOD provides an additional degree of freedom in this separation by also including considerations on the frequencies

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

involved. In the data-driven identification of coherent patterns, the mPOD allows overcoming the POD non-uniqueness problem in case of phenomena having similar energy content (e.g. Mendez *et al.* 2018*b*).

More generally, by providing bases that are spectrally cleaner than POD bases but energetically more relevant than DFT/DMD ones, the mPOD can be useful in any 'projection-based' application (from pattern recognition to low-order modelling of dynamical systems) in which the constraints of energy optimality or spectral purity are both unnecessarily extreme.

Acknowledgements

This research project has been partially funded by Arcelor-Mittal. The authors gratefully acknowledge the collaboration of the VUB-MECH laboratories at the Vrije Universiteit Brussel (VUB), for providing the time-resolved PIV system and the contribution of Y. Aksoy, who participated in the TR-PIV experiments in the framework of his Research Master program at the von Kármán Institute.

Supplementary movies

Supplementary movies are available at https://doi.org/10.1017/jfm.2019.212.

Appendix A. Notes on filter design and wavelet theory for MRA

The MRA presented in § 3.2 can operate with any digital filter, provided that the spectral overlapping between the transfer function of adjacent scales is reasonably small and the sum of the moduli of all the transfer functions in the filter bank yields unity. In the formulation of the mPOD presented in this work, we have used the simplest possible filters, namely finite impulse response (FIR) linear filters.

These are characterized by their impulse response $h[k]$, that is the inverse Fourier transform of the associated transfer function. The impulse response (or kernel) is the output of the filter when acting on a Dirac impulse $\delta[k]$, and its design can be carried out using standard methods presented in classical textbooks on signal analysis (Oppenheim & Schafer 2009; Hayes 2011). Among these, the current mPOD algorithm uses the windowing method, which consists in smoothing an ideal transfer function

$$\boldsymbol{H}^I_{\mathcal{L}m}(\omega) = \begin{cases} \mathrm{e}^{-\alpha\omega\mathrm{j}} & \text{for } |\omega_n| \leqslant 2\pi f^c_m, \\ 0 & \text{otherwise,} \end{cases} \tag{A 1}$$

with a kernel of size N, introducing a constant phase delay $\alpha = N/2$. The ideal transfer function for each scale is constructed from the frequency splitting vector $\boldsymbol{F}_V$. This leads to an unstable filter with infinite impulse response:

$$h^I_{\mathcal{L}m}[k] = \frac{1}{2\pi}\int_{-\pi}^{\pi} \boldsymbol{H}^I_{\mathcal{L}m}(\omega)\mathrm{e}^{-\mathrm{j}\omega k}\,\mathrm{d}\omega = \frac{\sin[2\pi f^c_m/f_s(k-\alpha)]}{\pi[k-\alpha]}, \tag{A 2}$$

but the smoothing in the frequency domain corresponds to a regularizing multiplication in the time domain:

$$h_{\mathcal{L}m}[k] = h^I_{\mathcal{L}m}[k]w[k]. \tag{A 3}$$

The linear phase response is provided by taking the window kernel to be symmetric about its midpoint, that is $w[n] = w[N-n]$. The choice of the window, and most

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

importantly its size N, controls how much the real transfer function differs from the ideal one, in terms of pass-band/stop-band deviation, and transition band (Oppenheim & Schafer 2009; Hayes 2011). Once the window is chosen, the low-pass transfer function corresponding to the impulse response in (A 3) is:

$$\boldsymbol{H}_{\mathcal{L}_m} = \frac{1}{2\pi} \int_{-\pi}^{\pi} \boldsymbol{H}^I_{\mathcal{L}m}(\omega)\widehat{w}(\omega - \theta)\, \mathrm{d}\theta, \tag{A 4}$$

with $\omega = 2\pi f_n/f_s$ the normalized frequency (radiant/samples) and $\widehat{w} = \boldsymbol{\Psi}_{\mathcal{F}} w$ the DFT of the window. Observe that (A 4) has an analytical representation and thus the transfer function of appropriate size $\boldsymbol{H}_{\mathcal{L}_m} \in \mathbb{C}^{1\times n_t}$, here constructed as a row vector for later convenience, can be computed from (A 4) on an arbitrary frequency discretization. In this work, the default window chosen is the Hanning window, which gives a stop-band attenuation of -44 dB (decibels) and transition width of $3.1/N$.

Concerning the filter order, one should observe that a sharp frequency cutoff becomes more and more difficult when lower frequencies are considered. This parameter is therefore left as a function of the frequency bandwidth considered, as it is common in the filter bank formulation of the continuous wavelet transform (Mallat 2009; Strang & Nguyen 1996). For the test case in § 6, for example, the filter orders N_m is

$$N_m = \begin{cases} 0.5n_t & \text{for } \boldsymbol{F}_V[m] \leqslant f_s/20, \\ n_t f_s/(40\boldsymbol{F}_V[m]) & \text{for } \boldsymbol{F}s/20 \leqslant \boldsymbol{F}_V[m] \leqslant f_s/2. \end{cases} \tag{A 5}$$

A zero-phase implementation of these filters, using for example the Matlab function filtfilt, is based on backward and forward filtering, in order to cancel the phase delay (Gustafsson 1996; Matlab 2017). This operation doubles the filter order, squares the modulus of its transfer function but renders it non-causal.

Finally, it is worth noticing that in previous formulations of the mPOD (see Mendez *et al.* 2018*b*,*a*), the filter bank architecture is presented in terms of discrete (dyadic) wavelet transform (DWT). This formulation replaces the concept of impulse response with that of wavelet (for high-pass) and scale (for low-pass) functions, constructed from dilatation and shifting of reference functions. The link between wavelet theory and filter banks is extensively discussed by Mallat (2009) and Strang & Nguyen (1996).

For the purposes of this work, it suffices to notice that the MRA of $\boldsymbol{K}$ can by carried out via wavelet toolboxes for image compression/denoising (Gonzalez & Woods 2007; Mallat 2009), based on the pyramid algorithm proposed by Mallat (1989). In this case, however, both the scaling and the shift of the reference wavelet/scale functions are fixed and the frequency splitting vector $\boldsymbol{F}_V$ becomes: $\boldsymbol{F}_V = [2^{-M}, \ldots, 2^{-3}, 2^{-2}, 2^{-1}]fs/2$ with f_s the sampling frequency and $M = \log_2(n_t)$ the number of scales.

REFERENCES

ANTORANZ, A., IANIRO, A., FLORES, O. & VILLALBA, M. G. 2018 Extended proper orthogonal decomposition of non-homogeneous thermal fields in a turbulent pipe flow. *Intl J. Heat Mass Transfer* **118**, 1264–1275.

AUBRY, N. 1991 On the hidden beauty of the proper orthogonal decomposition. *Theor. Comput. Fluid Dyn.* **2** (5–6), 339–352.

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

AUBRY, N., GUYONNET, R. & LIMA, R. 1991 Spatiotemporal analysis of complex signals: theory and applications. *J. Stat. Phys.* **64** (3–4), 683–739.

BARRERO-GIL, A., PINDADO, S. & AVILA, S. 2012 Extracting energy from vortex-induced vibrations: A parametric study. *Appl. Math. Model.* **36** (7), 3153–3160.

BERGMANN, M. & CORDIER, L. 2008 Optimal control of the cylinder wake in the laminar regime by trust-region methods and POD reduced-order models. *J. Comput. Phys.* **227** (16), 7813–7840.

BERKOOZ, G. 1992 Observations on the proper orthogonal decomposition. In *Studies in Turbulence*, pp. 229–247. Springer.

BERKOOZ, G., ELEZGARAY, J., HOLMES, P., LUMLEY, J. & POJE, A. 1994 The proper orthogonal decomposition, wavelets and modal approaches to the dynamics of coherent structures. *Appl. Sci. Res.* **53** (3–4), 321–338.

BERKOOZ, G., HOLMES, P. & LUMLEY, J. L. 1993 The proper orthogonal decomposition in the analysis of turbulent flows. *Annu. Rev. Fluid Mech.* **25** (1), 539–575.

BLANCHARD, R., NG, W., LOWE, T. K. & VANDSBURGER, U. 2014 Simulating bluff-body flameholders: on the use of proper orthogonal decomposition for wake dynamics validation. *Trans ASME J. Engng Gas Turbines Power* **136** (12), 122603.

BORÉE, J. 2003 Extended proper orthogonal decomposition: a tool to analyse correlated events in turbulent flows. *Exp. Fluids* **35** (2), 188–192.

BOURGEOIS, J. A., NOACK, B. R. & MARTINUZZI, R. J. 2013 Generalized phase average with applications to sensor-based flow estimation of the wall-mounted square cylinder wake. *J. Fluid Mech.* **736**, 316–350.

BROCKWELL, P. J. & DAVIS, R. A. 1987 *Time Series: Theory and Methods*. Springer.

BRUNTON, S. L., BRUNTON, B. W., PROCTOR, J. L. & KUTZ, J. N. 2016*a* Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PLOS ONE* **11** (2), e0150171.

BRUNTON, S. L. & NOACK, B. R. 2015 Closed-loop turbulence control: progress and challenges. *Appl. Mech. Rev.* **67** (5), 050801.

BRUNTON, S. L., PROCTOR, J. L. & KUTZ, J. N. 2016*b* Sparse identification of nonlinear dynamics with control (sindyc). *IFAC-Papers OnLine* **49** (18), 710–715.

BUCHLIN, J.-M. 2011 Convective heat transfer in impinging gas jet arrangements. *J. Appl. Fluid Mech.* **4** (2), 137–149.

BUDIŠIĆ, M., MOHR, R. & MEZIĆ, I. 2012 Applied koopmanism. *Chaos: An Interdisciplinary J. Nonlinear Sci.* **22** (4), 047510.

CAMMILLERI, A., GUENIAT, F., CARLIER, J., PASTUR, L., MEMIN, E., LUSSEYRAN, F. & ARTANA, G. 2013 POD-spectral decomposition for fluid flow analysis and model reduction. *Theor. Comput. Fluid Dyn.* **27** (6), 787–815.

DEL CASTILLO-NEGRETE, D., SPONG, D. A. & HIRSHMAN, S. P. 2008 Proper orthogonal decomposition methods for noise reduction in particle-based transport calculations. *Phys. Plasmas* **15** (9), 092308.

CHARMIYAN, M., AZIMIAN, A. R., SHIRANI, E., ALOUI, F., DEGOUET, C. & MICHAELIS, D. 2017 3D tomographic PIV, POD and vortex identification of turbulent slot jet flow impinging on a flat plate. *J. Mech. Sci. Technol.* **31** (11), 5347–5357.

CHARRU, F. & DE FORCRAND-MILLARD, P. 2009 *Hydrodynamic Instabilities*. Cambridge University Press.

CHATTERJEE, A. 2000 An introduction to the proper orthogonal decomposition. *Curr. Sci.* **78** (7), 808–817.

CHEN, K. K., TU, J. H. & ROWLEY, C. W. 2012 Variants of dynamic mode decomposition: boundary condition, koopman, and fourier analyses. *J. Nonlinear Sci.* **22** (6), 887–915.

CITRINITI, J. H. & GEORGE, W. K. 2000 Reconstruction of the global velocity field in the axisymmetric mixing layer utilizing the proper orthogonal decomposition. *J. Fluid Mech.* **418**, 137–166.

CORDIER, L. & BERGMANN, M. 2013 Proper orthogonal decomposition: an overview. In *Advanced Post-Processing of Experimental and Numerical Data* (ed. L. David & C. Schram), von Karman Institute for Fluid Dynamics, VKI – LS 2.

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

DAUBECHIES, I. 1992 *Ten Lectures on Wavelets*. Society for Industrial and Applied Mathematics.

DE SOUZA, F., DELVILLE, J., LEWALLE, J. & BONNET, J. P. 1999 Large scale coherent structures in a turbulent boundary layer interacting with a cylinder wake. *Exp. Therm. Fluid Sci.* **19**, 204–213.

DEANE, A. E., KEVREKIDIS, I. G., KARNIADAKIS, G. E. & ORSZAG, S. A. 1991 Low-dimensional models for complex geometry flows: application to grooved channels and circular cylinders. *Phys. Fluids (A/Fluid)* **3** (10), 2337–2354.

DIDDEN, N. & HO, C.-M. 1985 Unsteady separation in a boundary layer produced by an impinging jet. *J. Fluid Mech.* **160**, 235–256.

DURIEZ, T., BRUNTON, S. L. & NOACK, B. R. 2017 *Machine Learning Control – Taming Nonlinear Dynamics and Turbulence*. Springer International.

DUWIG, C. & IUDICIANI, P. 2009 Extended proper orthogonal decomposition for analysis of unsteady flames. *Flow Turbul. Combust.* **84** (1), 25–47.

ERICHSON, N. B., BRUNTON, S. L. & KUTZ, J. N. 2017 Randomized dynamic mode decomposition. arXiv:1702.02912v1.

FAHL, M. 2000 Trust-region methods for flow control based on reduced order modeling. PhD thesis, Université de Trier.

FARGE, M. 1992 Wavelet transforms and their applications to turbulence. *Annu. Rev. Fluid Mech.* **24** (1), 395–458.

FLÓR, J. B., VAN HEIJST, G. J. F. & DELFOS, R. 1995 Decay of dipolar vortex structures in a stratified fluid. *Phys. Fluids* **7** (2), 374–383.

FOLLAND, G. B. 2009 *Fourier Analysis and its Applications*, Pure and Applied Undergraduate Texts. American Mathematical Society.

GEORGE, W. K. 1988 Insight into the dynamics of coherent structures from a proper orthogonal decomposition. In *Near Wall Turbulence* (ed. S. Kline *et al.*), Hemisphere.

GLAUSER, M. N., LEIB, S. J. & GEORGE, W. K. 1987 Coherent structures in the axisymmetric turbulent jet mixing layer. In *Turbulent Shear Flows 5*, pp. 134–145. Springer.

GLEZER, A., KADIOGLU, Z. & PEARLSTEIN, A. J. 1989 Development of an extended proper orthogonal decomposition and its application to a time periodically forced plane mixing layer. *Phys. Fluids (A/Fluid)* **1** (8), 1363–1373.

GOLUB, G. H. 2013 *Matrix Computations*. J. Hopkins University Press.

GONZALEZ, R. C. & WOODS, R. E. 2007 *Digital Image Processing*, 3rd edn. Pearson.

GORDEYEV, S. V. & THOMAS, F. O. 2000 Coherent structure in the turbulent planar jet. Part 1. Extraction of proper orthogonal decomposition eigenmodes and their self-similarity. *J. Fluid Mech.* **414**, 145–194.

GRAY, R. M. 2005 Toeplitz and circulant matrices: a review. *Foundations and Trends® in Communications and Information Theory*, vol. 2(3), pp. 155–239. Now.

GRENANDER, U. 2001 *Toeplitz Forms and Their Applications*. American Mathematical Society.

GRINSTEIN, F. F. & DEVORE, C. R. 1996 Dynamics of coherent structures and transition to turbulence in free square jets. *Phys. Fluids* **8** (5), 1237–1251.

GROSEK, J. & KUTZ, J. N. 2014 Dynamic mode decomposition for real-time background/foreground separation in video. arXiv:1404.7592v1.

GUDMUNDSSON, K. & COLONIUS, T. 2011 Instability wave models for the near-field fluctuations of turbulent jets. *J. Fluid Mech.* **689**, 97–128.

GUSTAFSSON, F. 1996 Determining the initial states in forward-backward filtering. *IEEE Trans. Signal Process.* **44** (4), 988–992.

GUTMARK, E., WOLFSHTEIN, M. & WYGNANSKI, I. 1978 The plane turbulent impinging jet. *J. Fluid Mech.* **88** (04), 737.

HADŽIABDIĆ, M. & HANJALIĆ, K. 2008 Vortical structures and heat transfer in a round impinging jet. *J. Fluid Mech.* **596**, 221–260.

HAMMAD, K. J. & MILANOVIC, I. M. 2009 A POD study of an impinging jet flowfield. In *Volume 1: Symposia, Parts A, B and C*. ASME.

HARRIS, F. J. 1978 On the use of windows for harmonic analysis with the discrete fourier transform. *Proc. IEEE* **66** (1), 51–83.

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

HAYES, M. H. 2011 *Schaums Outline of Digital Signal Processing*. McGraw-Hill Education Ltd.

HOFFMANN, K. A. & CHIANG, S. T. 1993 *Computational Fluid Dynamics for Engineers*. Engineering Education System.

HOLMES, P., LUMLEY, J. L., BERKOOZ, G. & ROWLEY, C. 2012 *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, 2nd edn. Cambridge University Press.

HOLMES, P. J., LUMLEY, J. L., BERKOOZ, G., MATTINGLY, J. C. & WITTENBERG, R. W. 1997 Low-dimensional models of coherent structures in turbulence. *Phys. Rep.* **287** (4), 337–384.

HUSSAIN, A. K. M. F. 1986 Coherent structures and turbulence. *J. Fluid Mech.* **173**, 303–356.

JACKSON, J. E. 1991 *A User's Guide to Principal Components*. Wiley-Interscience.

JOVANOVIĆ, M. R., SCHMID, P. J. & NICHOLS, J. W. 2014 Sparsity-promoting dynamic mode decomposition. *Phys. Fluids* **26** (2), 024103.

KAISER, G. 2010 *A Friendly Guide to Wavelets*. Springer.

KIM, K. C., MIN, Y. U., OH, S. J., AN, N. H., SEOUDI, B., CHUN, H. H. & LEE, I. 2007 Time-resolved PIV investigation on the unsteadiness of a low Reynolds number confined impinging jet. *J. Vis.* **10** (4), 367–379.

KOOPMAN, B. O. 1931 Hamiltonian systems and transformation in hilbert space. *Proc. Natl Acad. Sci.* **17** (5), 315–318.

KUHN, S., KENJEREŠ, S. & VON ROHR, P. R. 2010 Large eddy simulations of wall heat transfer and coherent structures in mixed convection over a wavy wall. *Intl J. Therm. Sci.* **49** (7), 1209–1226.

KUNISCH, K. & VOLKWEIN, S. 1999 Control of the burgers equation by a reduced-order approach using proper orthogonal decomposition. *J. Optim. Theor. Applics.* **102** (2), 345–371.

KUTZ, J. N. 2013 *Data-Driven Modeling & Scientific Computation*. Oxford University Press.

KUTZ, J. N., BRUNTON, S. L., BRUNTON, B. W. & PROCTOR, J. L. 2016*a* *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. SIAM-Society for Industrial and Applied Mathematics.

KUTZ, J. N., FU, X. & BRUNTON, S. L. 2016*b* Multiresolution dynamic mode decomposition. *SIAM J. Appl. Dyn. Syst.* **15** (2), 713–735.

LOÈVE, M. 1977 *Probability Theory I*. Springer.

LOISEAU, J.-C., NOACK, B. R. & BRUNTON, S. L. 2018 Sparse reduced-order modelling: sensor-based dynamics to full-state estimation. *J. Fluid Mech.* **844**, 459–490.

LUMLEY, J. L. 1967 The structure of inhomogeneous turbulent flows. In *Atmospheric Turbulence and Radio Propagation* (ed. A. M. Yaglom & V. I. Tatarski), pp. 166–178. Nauka.

LUMLEY, J. L. 1970 *Stochastic Tools in Turbulence*. Dover.

MALLAT, S. G. 1989 A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans. Pattern Anal. Mach. Intell.* **11** (7), 674–693.

MALLAT, S. G. 2009 *A Wavelet Tour of Signal Processing*. Elsevier.

MALLOR, F., VILA, C. S., IANIRO, A. & DISCETTI, S. 2018 Wall-mounted perforated cubes in a boundary layer: local heat transfer enhancement and control. *Intl J. Heat Mass Transfer* **117**, 498–507.

MARIOTTI, A., LEGRAS, B. & DRITSCHEL, D. G. 1994 Vortex stripping and the erosion of coherent structures in two-dimensional flows. *Phys. Fluids* **6** (12), 3954–3962.

MATLAB 2017 *Signal Processing Toolbox*. The MathWorks.

MAUREL, S., BORÉE, J. & LUMLEY, J. L. 2001 Extended proper orthogonal decomposition: application to jet/vortex interaction. *Flow Turbul. Combust.* **67** (2), 125–136.

MELNIKOV, K., KREILOS, T. & ECKHARDT, B. 2014 Long-wavelength instability of coherent structures in plane couette flow. *Phys. Rev.* E **89**, 043008.

MENDEZ, M. A., BALABANE, M. & BUCHLIN, J.-M. 2018*a* Multi-scale proper orthogonal decomposition (mPOD). In *AIP Conference Proceedings*. vol. **1978**, 060018, AIP.

MENDEZ, M. A., RAIOLA, M., MASULLO, A., DISCETTI, S., IANIRO, A., THEUNISSEN, R. & BUCHLIN, J.-M. 2017 POD-based background removal for particle image velocimetry. *Exp. Therm. Fluid Sci.* **80**, 181–192.

MENDEZ, M. A., SCELZO, M. T. & BUCHLIN, J.-M. 2018*b* Multiscale modal analysis of an oscillating impinging gas jet. *Exp. Therm. Fluid Sci.* **91**, 256–276.

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

MEZIĆ, I. 2005 Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dyn.* **41** (1–3), 309–325.

MEZIĆ, I. 2013 Analysis of fluid flows via spectral properties of the Koopman operator. *Annu. Rev. Fluid Mech.* **45** (1), 357–378.

MIRANDA, A. A., BORGNE, Y.-A. & LE BONTEMPI, G. 2008 New routes from minimal approximation error to principal components. *Neural Proc. Lett.* **27** (3), 197–207.

MLADIN, E.-C. & ZUMBRUNNEN, D. A. 2000 Alterations to coherent flow structures and heat transfer due to pulsations in an impinging air-jet. *Intl J. Therm. Sci.* **39** (2), 236–248.

MUMFORD, J. C. 1982 The structure of the large eddies in fully developed turbulent shear flows. Part 1. The plane jet. *J. Fluid Mech.* **118**, 241–268.

NAGARAJAN, K. K., SINGHA, S., CORDIER, L. & AIRIAU, C. 2018 Open-loop control of cavity noise using proper orthogonal decomposition reduced-order model. *Comput. Fluids* **160**, 1–13.

NOACK, B. R. 2016 From snapshots to modal expansions – bridging low residuals and pure frequencies. *J. Fluid Mech.* **802**, 1–4.

NOACK, B. R., AFANASIEV, K., MORZYŃSKI, M., TADMOR, G. & THIELE, F. 2003 A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *J. Fluid Mech.* **497**, 335–363.

NOACK, B. R., STANKIEWICZ, W., MORZYŃSKI, M. & SCHMID, P. J. 2016 Recursive dynamic mode decomposition of transient and post-transient wake flows. *J. Fluid Mech.* **809**, 843–872.

OLIVER, N. M., ROSARIO, B. & PENTLAND, A. P. 2000 A bayesian computer vision system for modeling human interactions. *IEEE Trans. Pattern Anal. Mach. Intell.* **22** (8), 831–843.

OPPENHEIM, A. V. & SCHAFER, R. W. 2009 *Discrete-Time Signal Processing [With Access Code]*. Addison–Wesley.

PICARD, C. & DELVILLE, J. 2000 Pressure velocity coupling in a subsonic round jet. *Intl J. Heat Fluid Flow* **21** (3), 359–364.

PIERIS, S., ZHANG, X., YARUSEVYCH, S. & PETERSON, S. D. 2017 Evolution of coherent structures in a two-dimensional impinging jet. In *10th International Symposium on Turbulence and Shear Flow Phenomena (TSFP10), Chicago, USA*, pp. 89–94.

POLLARD, A., CASTILLO, L., DANAILA, L. & GLAUSER, M. (Eds) 2017 *Whither Turbulence and Big Data in the 21st Century?* Springer International.

RAHMANIAN, M., CHENG, L., ZHAO, M. & ZHOU, T. 2014 Vortex induced vibration and vortex shedding characteristics of two side-by-side circular cylinders of different diameters in close proximity in steady flow. *J. Fluids Struct.* **48**, 260–279.

RINOSHIKA, A. & ZHOU, Y. 2005 Orthogonal wavelet multi-resolution analysis of a turbulent cylinder wake. *J. Fluid Mech.* **524**, 229–248.

ROWLEY, C. W. & DAWSON, S. T. M. 2017 Model reduction for flow analysis and control. *Annu. Rev. Fluid Mech.* **49** (1), 387–417.

ROWLEY, C. W., MEZIĆ, I., BAGHERI, S., SCHLATTER, P. & HENNINGSON, D. S. 2009 Spectral analysis of nonlinear flows. *J. Fluid Mech.* **641**, 115–127.

SCHMID, P. J. 2010 Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* **656**, 5–28.

SCHMID, P. J. & BRANDT, L. 2014 Analysis of fluid systems: stability, receptivity, Sensitivity Lecture notes from the FLOW-NORDITA summer school on advanced instability methods for complex flows, stockholm, sweden, 2013. *Appl. Mech. Rev.* **66** (2), 021003.

SCHNEIDER, K. & VASILYEV, O. V. 2010 Wavelet methods in computational fluid dynamics. *Annu. Rev. Fluid Mech.* **42** (1), 473–503.

SCHRIJER, F. F. J., SCIACCHITANO, A. & SCARANO, F. 2014 Spatio-temporal and modal analysis of unsteady fluctuations in a high-subsonic base flow. *Phys. Fluids* **26** (8), 086101.

SHUKLA, K. K. & TIWARI, A. K. 2013 *Efficient Algorithms for Discrete Wavelet Transform*. Springer.

SIEBER, M., PASCHEREIT, C. O. & OBERLEITHNER, K. 2016 Spectral proper orthogonal decomposition. *J. Fluid Mech.* **792**, 798–828.

SIEBER, M., PASCHEREIT, C. O. & OBERLEITHNER, K. 2017 On the nature of spectral proper orthogonal decomposition and related modal decompositions. arXiv:1712.08054v1.

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press

SIROVICH, L. 1987 Turbulence and the dynamics of coherent structures. Part I. Coherent structures. *Q. Appl. Maths* **45** (3), 561–571.

SIROVICH, L. 1989 Chaotic dynamics of coherent structures. *Physica* D **37** (1), 126–145.

SIROVICH, L. 1991 Analysis of turbulent flows by means of the empirical eigenfunctions. *Fluid Dyn. Res.* **8**, 85–100.

SOBRAL, A., JAVED, S., JUNG, S. K., BOUWMANS, T. & ZAHZAH, E. 2015 Online stochastic tensor decomposition for background subtraction in multispectral video sequences. In *2015 IEEE International Conference on Computer Vision Workshop (ICCVW)*, IEEE.

SPECHT, E. 2014 Impinging jet drying. In *Modern Drying Technology, Process Intensification* (ed. E. Tsotsas & A. S. Mujumdar), vol. 5, chap. 1, pp. 1–26. Wiley-VCH.

STEWART, G. W. 1993 On the early history of the singular value decomposition. *SIAM Rev.* **35** (4), 551–566.

STRANG, G. & NGUYEN, T. 1996 *Wavelets and Filter Banks*. Wellesley-Cambridge Press.

SURESH, P. R., SRINIVASAN, K., SUNDARARAJAN, T. & DAS, S. K. 2008 Reynolds number dependence of plane jet development in the transitional regime. *Phys. Fluids* **20** (4), 044105.

TAIRA, K., BRUNTON, S. L., DAWSON, S. T. M., ROWLEY, C. W., COLONIUS, T., MCKEON, B. J., SCHMIDT, O. T., GORDEYEV, S., THEOFILIS, V. & UKEILEY, L. S. 2017 Modal analysis of fluid flows: an overview. *AIAA J.* **55** (12), 4013–4041.

THIELICKE, W. & STAMHUIS, E. J. 2014 PIVlab – towards user-friendly, affordable and accurate digital particle image velocimetry in MATLAB. *J. Open Res. Softw.* **2**, e30.

THOMAS, F. O. & GOLDSCHMIDT, V. W. 1986 Structural characteristics of a developing turbulent planar jet. *J. Fluid Mech.* **163** (1), 227.

TOWNE, A., SCHMIDT, O. T. & COLONIUS, T. 2018 Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *J. Fluid Mech.* **847**, 821–867.

TRIELING, R. R., VAN WSENBEECK, J. M. A. & VAN HEIJST, G. J. F. 1998 Dipolar vortices in a strain flow. *Phys. Fluids* **10** (1), 144–159.

TU, J., ROWLEY, C. W., LUCHTENBURG, D. M., BRUNTON, S. L. & KUTZ, J. N. 2014 On dynamic mode decomposition: theory and applications. *J. Comput. Dyn.* **1** (2), 391–421.

VAIDYANATHANM, P. P. 1992 *Multirate Systems and Filter Banks*. Prentice Hall.

VENTURI, D. & KARNIADAKIS, G. E. M. 2004 Gappy data and reconstruction procedures for flow past a cylinder. *J. Fluid Mech.* **519**, 315–336.

VOLKWEIN, S. 2013 *Proper Orthogonal Decomposition: Theory and Reduced-Order Modelling*. University of Konstanz, Department of Mathematics and Statistics.

WELCH, P. 1967 The use of fast fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms. *IEEE Trans. Audio Electroacoust.* **15** (2), 70–73.

WILLCOX, K. 2006 Unsteady flow sensing and estimation via the gappy proper orthogonal decomposition. *Comput. Fluids* **35** (2), 208–226.

ZUCKERMAN, N. & LIOR, N. 2006 Jet impingement heat transfer: physics, correlations, and numerical modeling. *Adv. Heat Transfer* **39**, 565–631.

https://doi.org/10.1017/jfm.2019.212 Published online by Cambridge University Press